A Role-Based Access Control Model and Reference Implementation Within a Corporate Intranet

DAVID F. FERRAIOLO, JOHN F. BARKLEY, and D. RICHARD KUHN
National Institute of Standards and Technology

This paper describes NIST's enhanced RBAC model and our approach to designing and implementing RBAC features for networked Web servers. The RBAC model formalized in this paper is based on the properties that were first described in Ferraiolo and Kuhn [1992] and Ferraiolo et al. [1995], with adjustments resulting from experience gained by prototype implementations, market analysis, and observations made by Jansen [1988] and Hoffman [1996]. The implementation of RBAC for the web (RBAC/Web) provides an alternative to the conventional means of administering and enforcing authorization policy on a server-by-server basis. RBAC/Web provides administrators with a means of managing authorization data at the enterprise level, in a manner consistent with the current set of laws, regulations, and practices.

1. INTRODUCTION

In recent years, vendors have begun implementing role-based access control (RBAC) features into their databases, system management, and operating system products, without any general agreement as to what actually constitutes an appropriate set of RBAC features. Several RBAC models were proposed [Ferraiolo and Kuhn 1992; Ferraiolo et al. 1995; Sandhu et al. 1996; Nyanchama and Osborn 1994], without any attempt at standardizing salient RBAC features. To identify RBAC features that both exhibit true enterprise value and are practical to implement, the National Insti-

tute of Standards and Technology (NIST) has conducted and sponsored market analysis [Ferraiolo et al. 1993], developed prototype implementations, sponsored external research [Feinstein 1995], and published formal RBAC models [Ferraiolo and Kuhn 1992; Ferraiolo et al. 1995]. As a result of these and other efforts, much has been learned about RBAC and its practical implementation. The purpose of this paper is to rigorously define NIST's enhanced RBAC model and to describe its reference implementation within a world wide web (WWW) intranet application (RBAC/Web). An intranet application was chosen to show RBAC's application to a network and to demonstrate its ability to solve common authorization management and policy problems.

1.1 The Existing Problem

One of the greatest obstacles to the growth of intranets as a means of enterprise computing is inability to manage authorization data effectively. Authorization management today is costly and prone to error. Web server administrators usually control user access to enterprise published documents through creation and maintenance of access control lists (ACLs) on a server-by-server basis. ACLs specify, for each protected resource, a list of named individuals, or groups composed of individual users, with their respective modes of access to that object.

This use of ACLs is problematic for a variety of reasons. ACLs are tied to particular objects. As such, they are appropriate for discretionary need-to-know policies, where ownership of objects resides with the end user. In many enterprises within industry and civilian government, end users do not "own" the information to which they are allowed access [Ferraiolo et al. 1993; Ferraiolo and Kuhn 1992]. For these organizations, the corporation or agency is the actual "owner" of system objects and discretionary control on the part of the users may not be appropriate. Although enforcing a need-to-know policy is important where classified information is of concern, there exists a general need to support subject-based security policies, such as access based on competency, the enforcement of conflict-of-interest rules, or permitting access based on a strict concept of least privilege. To support such policies requires the ability to restrict access based on a user function or role within the enterprise. Here the relevant question is: What are the current operational capabilities for this user?

ACLs further complicate matters by allowing the direct association of users with permissions. A large number of users, each with many permissions, implies a very large number of user/permission associations that have to be managed. Thus, when a user takes on different responsibilities within the enterprise, reflecting these changes entails a thorough review, resulting in the selective addition or deletion of user/permission associations on all servers. The larger the number of user/permission associations to be managed, the greater the risk of maintaining residual and inappropriate user access rights.

Due to the potential problems associated with this lack of operational security assurance, organizations have resisted publishing sensitive infor-

mation on their web servers. This limits their utility, depriving the organization of potential productivity gains.

1.2 Solution: Role-Based Access Control

Although it can support discretionary access control (DAC) policies[Sandhu and Munawer 1998], RBAC is primarily a nondiscretionary access control model. RBAC does not permit users to be directly associated with permissions.

   With RBAC, permissions are authorized for roles and roles are authorized for users. Thus, when administering RBAC, two different types of associations must be managed: associations between users and roles and associations between roles and permissions. When a user's job position changes, only the user/role associations change. If the job position is represented by a single role, then, when a user's job position changes, there are only two user/role associations to change: remove the association between the user and the user's current role and add an association between the user and the user's new role.

   1.2.1 *Reduced Administrative Cost Complexity*.   There is usually a direct relationship between the cost of administration and the number of associations that must be managed to administer an access control policy: the larger the number of associations, the costlier and more error-prone the access control administration. In most organizations, the use of RBAC reduces the number of associations that must be managed.

   Job positions are typically occupied by more than one individual and most positions require more than one permission in order for an individual in a job position to carry out the responsibilities of that position. One can describe the associations authorizing permissions to individuals who perform the responsibilities of a job position as an ordered pair consisting of a set of individuals and a set of permissions, that is,

$$(U, P)$$

where U = the set of individuals in a job position and P = the set of permissions required to perform that job position.

   The number of associations required to directly relate the individuals to those permissions is

$$|U| \cdot |P|$$

where $|U|$ = the number of individuals in the set U and $|P|$ = the number of permissions in the set P.

   A role can be described as a set of permissions. Thus, the set P can refer to a role that is the job position whose user/role and role/permission associations are represented by the ordered pair (U, P). The number of user/role and role/permission associations required to authorize each user

in the set U for each of the permissions in the set P where P represents a role is

$$|U| + |P|.$$

That is, there is an association with the role P for each individual in U and an association with the role P for each permission in P.

For a job position, if $|U| + |P| < |U| \cdot |P|$, then the administrative advantage of RBAC over relating users directly with permissions is realized for that job position. A sufficient condition for $|U| + |P| < |U| \cdot |P|$ is $|U|, |P| > 2$, which is typically the case for most job positions in most organizations.

If n_{jp} is the number of job positions within an organization, then the administrative advantage of RBAC is realized organization-wide when .

$$\Sigma_i^{n_{jp}}(|U_i| + |P_i|) < \Sigma_i^{n_{jp}}(|U_i| \cdot |P_i|).$$

1.2.2 *Policy Support.* The limiting factor in effectively enforcing security policy is not the capability of the ACL mechanism but the administrative interface. RBAC provides administrators with a context for the specification and enforcement of complex security policies that are often impractical to enforce through the direct administration of lower level access control mechanisms, such as ACLs.

Unlike ACLs that only support the specification of user/permission and group/permission relationships, the RBAC model supports the specification of user/role and role/role relationships. In particular, the RBAC model supports the specification of

—user/role associations, i.e., user authorizations to perform roles;

—role hierarchies, e.g., the role bank teller inherits all of the permissions of the role bank employee;

—separation of duty constraints (role/role associations indicating conflict of interest):
  static (SSD): a user cannot be authorized for both roles, e.g., bank teller and auditor;
  dynamic (DSD): a user can be authorized for both roles but cannot act simultaneously in both roles, e.g., a bank teller who has an account in the bank where employed cannot act in the role *teller* and *customer* simultaneously;

—limits on the number of users that can be authorized for a role (role cardinality), e.g., a branch office of a bank has only one branch manager.

For web server applications, RBAC provides administrative conveniences by composing the seemingly unrelated and incomprehensible authorization data of the lower level access control mechanisms, and other RBAC relational data, into a single RBAC authorization database. In doing so,

RBAC/Web, the implementation of the NIST RBAC model for intranets, organizes this authorization data and presents it to the intranet administrator(s) in a relational format and at a level of abstraction that is natural to the way enterprises are normally structured and conduct business. From an administrator's perspective, RBAC/Web serves as a visualization and maintenance tool[1] of the enterprise's intranet access control policies in terms of its users, roles, role hierarchies, operational constraints, and permissions.

Even with the enhanced administrative interface, there exist opportunities for inconsistencies that could result in undesirable security consequences. The complexity of dealing with consistency issues is not delegated to the administrator, but instead is handled by the RBAC/Web administrative software through the implementation of a series of integrity checks that are derived from the RBAC model [Gavrila and Barkley 1998]. Based on RBAC relationships, an interface can be built to solve virtually any authorization problem— a proposition demonstrated by the advanced security policies described in Simon and Zurko [1997] and Sandhu et al. [1996]. Still other real world policies can be expressed by imposing new constraints on time and location. The scope of the RBAC model formalized in this paper has been limited to those properties that were first described in Ferraiolo and Kuhn [1992], with adjustments that have resulted through observations made by Jansen [1988] and Hoffman [1996].

The remainder of this paper describes NIST's enhanced RBAC model and our approach to RBAC/Web. Section 2 provides a detailed description of the RBAC model with exemplary real-world security policies supported by the model. Section 3 provides an overview of the administrative approach assumed by the RBAC model and a description of our approach for implementing the RBAC model presented in Section 2. This includes RBAC/Web constituent process components and services, description of the RBAC/Web distributed authorization database, and a comprehensive scenario of use, from a client request for URL access (at the browser), through user authorization and role activation, to the result provided by the web server.

2. THE RBAC MODEL AND SUPPORTED POLICIES

The RBAC security model is both abstract and general. It is abstract because properties not relevant to security are not included; it is general because many designs could be valid interpretations of the model. The model allows design decisions to be postponed and is usable as a basis for the design of a variety of IT systems. A goal in creating the model was to provide as concise and usable a notation as possible so that the security-relevant properties of the model are not obscured by excessive notational detail.

---

[1]The RBAC/Web implementation for Windows NT provides a tool, RGP-Admin, to manage role-permission relationships (see Section 3.4), but in the UNIX implementation, the relationships are managed using the tools of the web server.

The RBAC model described below is sufficient to support a variety of security policies. In particular, an argument is made for least privilege and separation of duty. Least privilege is the time-honored administrative practice of selectively assigning privileges to users such that the user is given no more privilege than is necessary to perform his/her job. The principle of least privilege avoids the problem of an individual with the ability to perform unnecessary and potentially harmful actions as merely a side-effect of granting the ability to perform desired functions. Permissions (or privileges) are rights granted to an individual, or a subject acting on behalf of a user, that enable the holder of those rights to act in the system within the bounds of those rights. The question then becomes how to assign the set of system privileges to the aggregates of functions or duties that correspond to a role of a user or subject acting on behalf of the user. Least privilege provides a rationale for where to install the separation boundaries that are to be provided by RBAC protection mechanisms. Ensuring adherence to the principle of least privilege is largely an administrative challenge that requires identification of job functions, specification of the set of privileges required to perform each function, and restricton of the user to a domain with those privileges and nothing more.

Separation of duty refers to the partitioning of tasks and associated privileges among different roles associated with a single user to prevent users from colluding with one another. These separation concepts include multiplexing of shared resources, naming distinctive sets of permissions to include functional decomposition, categorical classification, and hierarchical decomposition. Within an RBAC system, these separation concepts are supported by the establishment of the principle of least privilege.

The RBAC security model has two components MC_0, and MC_1. Model component MC_0, called the *RBAC Authorization Database model*, defines the RBAC security properties for the authorization into static roles. Model component MC_1, called the *RBAC Activation model*, defines the RBAC security properties for dynamic activation of roles. Component MC_1 requires MC_0. Each model component is defined by the following subcomponents:

—a set of *types*, which define a basic set of elements, together with functions on these elements. There are two types of functions: mapping functions, which represent relationships between elements, and constraint functions, which represent restrictions on relationships between elements;

—a set of *rules*, which represent assumptions about the model; and

—a set of *properties*, which are implications of the rules.

2.1 MC_0: RBAC Authorization Database

The MC_0 model is defined in terms of mapping functions and static properties. Static properties refer to properties of the model that do not

involve either subject or mappings from subject to other basic elements (e.g., subject-user and active-roles). Static properties represent the most fundamental and the strongest constraints and relationships in the model. They include *role hierarchy*, *inheritance*, *cardinality*, and *static separation of duty*.

2.2 Users, Roles, and Permissions

RBAC is described in terms of individual users associated with roles as well as roles associated with permissions. As such, a role is a means for naming many-to-many relationships among individual users and unique permissions. A user in this model is a human being. A role is a job function or job title within the organization with some associated semantics regarding the authority and responsibility conferred on the user authorized for the role. We discuss role association with permissions below, but first we formally model user association with a role by the following types and mapping function:

type *user* of individual users;

type *role* of role identifiers;

$RM(r : role) \rightarrow 2^{user}$, the role/members mapping, which gives the set of users authorized for the role, r.

A role is called *empty* if and only if $RM[r] = \emptyset$.

A permission or privilege is an approval of a particular operation to be performed on one or more objects. The relationship between roles and permissions is depicted in Figure 1. As shown in Figure 1, arrows indicate a many-to-many relationship (i.e., a permission can be associated with one or more roles and a role can be associated with one or more permissions). We formally describe permission and the association of permission with a role by the following type and state independent mapping function:

type $permission = 2^{(operation \times object)}$; and

$RP(r : role) \rightarrow 2^{permissions}$, the role/permissions mapping, which gives the set of permissions authorized for the role r.

An operation is an executable image of a program, which, upon invocation, causes information to flow from or to one or more RBAC-protected objects, or cause the consumption of an exhaustible system resource. The type of operations and the objects that RBAC controls depend on the type of system in which it will be implemented. For example, within an operating system, operations might include read, write, and execute; within a database management system, operations might include insert, delete, select, append, and update. Operations are defined using the following type:

type *operation* of executable program.

Fig. 1. Role-operation-object relationship.

The purpose of any access control mechanism is to protect information and other resources. However, in applying RBAC to a computer system, we speak of protected objects. For a system that implements RBAC, the objects can represent information containers (e.g., files, directories in an operating system, or columns, rows, tables, and views within a database management system) or objects can represent exhaustible system resources such as printers, disk space, and CPU cycles. The set of objects covered by RBAC includes all of the objects included in the permissions that are associated with roles. For convenience, we introduce the following supertype, of which information-container and exhaustible-resource are disjoint subtypes:

$$\text{type } object = InformationContainer \cup ExhaustibleResource.$$

The following state-independent mapping functions are defined for permissions:

$POp(p : permission) \rightarrow \{operation\}$, the permission to operation mapping, which gives the set of operations associated with permission p.

$POb(p : permission) \rightarrow \{object\}$, the permission to object mapping, which gives the set of objects associated with permission p.

Figure 1 gives a relational representation of the set of RBAC elements originally formalized in Ferraiolo et al. [1995]. The arrows are used to represent many-to-many relationships (e.g., user-role, and role-operations, role-object). The lines represent relationships derived from the model component, MC_0, and as such are representative of the static relationships among the types as they pertain to the RBAC authorization database.
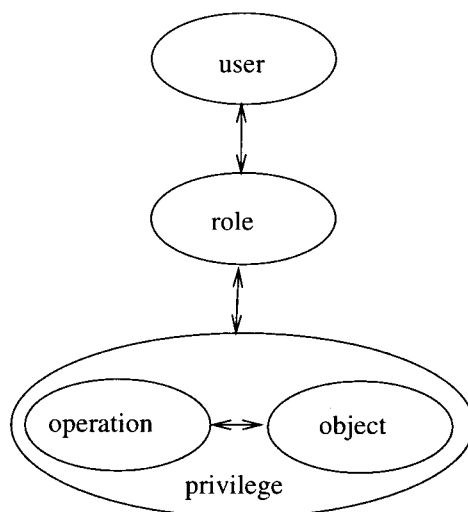
Fig. 2. Model element relationships.

For notational and conceptual purposes, the ternary relationship among role, operation, and object is dropped and is reformulated as a pair of binary relations, illustrated in Figure 2: one between operation and object, referred to as permission; the other between role and permission. Permission is a construct often associated with information systems, used to control user actions [Ferraiolo and Kuhn 1992]. In some contexts, permissions are also referred to as privileges.

This arrangement provides great flexibility and granularity in assigning privileges to roles and users to roles. Prior to providing these conveniences, it was often the case that a user was granted more access to resources than was needed due to the limited control over the type of access that can be associated with users and resources. Users may need to list directories and modify existing files, for example, without creating new files, or they may need to append records to a file without modifying existing records. Any increase in the flexibility of controlling access to resources also strengthens the application of the principle of least privilege. We now describe relationships between roles.

2.3 Role Membership Hierarchies

An instance of a role represents a many-to-many relationship between individual user members and individual permissions. To address policy and administrative authorization issues, RBAC includes the concept of containment. Containment is similar to inheritance in object-oriented systems, whereby the properties and constraints of a containing role are an extension of the properties and constraints of any contained role. Containment is also recursive; one role can contain other roles, which contain others, and so on. Besides facilitating role administration, containment permits the

substitution of role instances. That is, if role A contains role B, written $A \geq B$, then instances of role A are treated as instances of role B for the purpose of access control. Note that this rule defines inheritance properties for role membership, but not necessarily for role activation (see Section 2.6). Containment is defined by the following property:

Role membership inheritance. A containing role is effectively an instance of its contained roles:

$$(\forall i, j : role)(\forall u : user)i \geq j \wedge u \in RM[i] \Rightarrow u \in RM[j]. \qquad (1)$$

The complete set of containment relationships among all roles is referred to as a role hierarchy. The containment relation is represented as a set of ordered pairs $i \geq j$, where i is the parent or containing role, and j the child or contained role. We may also say that role i *inherits* role j, or that i is a *senior* and j is a *junior* role. The symbol " $\geq$ " denotes that the contains relationship forms a partial ordering of the roles (i.e., a reflexive, transitive, and antisymmetric relation). Thus, $i \geq j$ means i contains j, or alternatively, j is contained by i.

RBAC relations consist of an *explicit* set of role relationships (e.g., user-role, role-permission) and constraints (e.g., static separation of duty), and an *implicit* set of inherited relationships and constraints. The explicit relationships and constraints for a role are those relationships and constraints that are directly specified through an administrator role. The implicit relationships and constraints of a role are those relationships and constraints that are inherited from other roles within a hierarchy through containment. As such, inheritance is described in some models in terms of permissions [Sandhu et al. 1996;, Nyanchama and Osborn 1994]. For example, in Nyanchama and Osborn [1994], a consistent notion of containment is defined in terms of an is-junior relationship, where role, r_i is-junior to r_j, if the privileges (permissions) of r_i are a subset of the privileges of r_j. Implementing a permission inheritance scheme assumes a means of identifying and naming individual permissions, and directly and indirectly associating these permissions with each system role. For a self-contained RBAC system, such as a relational database management system, permissions can be centrally identified for any role. For systems that are implemented as a distributed RBAC authorization database, such as that of the RBAC/Web, the associated permissions for each role are not centrally available. For these systems, containment can be expressed in terms of inheritance of user-role relationships. The users authorized for a role include all users authorized for all roles contained by the role.

The inherited user members of a role include all user members of all roles that are contained by the role. For example, if u_1 is explicitly authorized for role A and role B is contained by role A, then u_1 is implicitly authorized for role B and can execute the permissions associated with role

B. In this sense the permissions of role B are inherited by role A. As such, RBAC/Web provides services for management of user authorization to roles in terms of user-role relationships, role-role relationships, and constraints on user-role relationships, while role-permission relationships are maintained on the back-end web servers. Because the NIST model places no restrictions on assignment of permissions to roles, such assignments with respect to sensitive objects must be under strict administrator control. This is similar to the nondiscretionary practice of assigning security levels to classified information in a DoD multilevel secure (MLS) environment.

As illustrated in Figure 3, role hierarchies can be combined to form a directed acyclic graph. A directed acyclic graph (DAG) role representation is a natural way of organizing roles to reflect authority, responsibility, and competency. The more powerful roles (i.e., those that contain the greatest number of permissions and are authorized for the fewest users) are found at the top of the graph, and the more general roles are at the bottom. Consider the Staff role; permissions of the Staff role are implicitly associated with all other roles in the enterprise, and the user members of all other roles are implicitly members of the Staff role. As shown in Figure 3, not all roles have to be related. For example, the roles AR Supervisor and Cashier Supervisor are not hierarchically related, and as such are not comparable under the contains relationship, but they may contain some or all of the same roles. It should also be noted that not all roles are intended to have explicit user members, and may exist for the sole purpose of associating permissions.

Role hierarchies provide a powerful and convenient means of administratively specifying and ensuring adherence to the principle of *least privilege* [Sandhu et al. 1996; Feinstein 1995]. Since many of the responsibilities overlap job categories, maximum privilege for each job category could cause unlawful access. To address least privilege issues, roles representing organizational functions can be associated with only those privileges that need to perform that function, and users can be made members of the roles that are known to be authorized to perform that function. In the cases where privileges completely overlap among roles, hierarchies of roles can be established. For example, it may seem sufficient to allow physicians to have access to all data within a patient record if their access is monitored sufficiently. However, this would entail much more auditing and monitoring than would be necessary with a better-defined access control mechanism. With RBAC, constraints can be placed on physician access so that, for example, only those fields that are associated with a particular type of physician can be accessed. Note that multiple inheritance in role hierarchies is allowed, but not required, by the model.

2.4 Static Constraints

RBAC provides administrators with the capability of imposing constraints on user-role and role-permission authorization. From a policy perspective, constraints provide a powerful means of enforcing conflicts of interest and

Fig. 3. Example role hierarchy.

cardinality rules for roles as they uniquely apply to an enterprise. The association of a user with a role can be subject to the following constraints:

—the role for which the user is authorized is not mutually exclusive with another role for which the user already possesses membership; and

—the numerical limitation that exists for the number of users authorized for a role cannot be exceeded.

2.4.1 *Static Separation of Duty*.   The first constraint, called *static separation of duty (SSD)*, can be used to enforce conflicts of interest policies that may arise as a result of a user gaining authorization for permissions associated with conflicting roles. This means that if a user is authorized as a member of one role, the user is prohibited from being a member of a second role. For example, a user who is authorized for the role Cashier in the application depicted in Figure 3 may not be authorized for the role Accounts Receivable Clerk. That is, the roles Cashier and Accounts Receivable Clerk are mutually exclusive. The SSD policy can be centrally specified and then be uniformly imposed on specific roles. The constraint function for mutually exclusive roles and the Static Separation of Duty property is specified as follows:

$Ea : role \times role$, the exclusive authorization set, which gives the pairs of roles that are mutually exclusive for role membership.

Static separation of duty. A user is authorized for a role only if that role is not mutually exclusive with any of the other roles for which the user is already authorized:

$$(\forall u : user)(\forall i, j : roles)u \in RM[i] \wedge u \in RM[j] \Rightarrow (i, j) \notin Ea. \qquad (2)$$

As described earlier, constraints are inherited within a role hierarchy. For example, in Figure 3, since roles Accounts Receivable Clerk and Billing Clerk have an SSD relationship, AR Supervisor also has an SSD relationship with Billing Clerk. Another way of thinking about this is that any instance of AR Supervisor can be treated as an instance of Accounts Receivable Clerk. Therefore, the SSD constraint that Billing Clerk has with Accounts Receivable Clerk must also apply to AR Supervisor. Rather than including constraint inheritance as a formal property of role hierarchy, constraints are better treated as side effects of imposing constraints on roles within a hierarchy.

Because a containing role is effectively an instance of its contained roles, no SSD relationship can exist between them. In the previous example, it would not make sense to have an SSD relationship between AR Supervisor and AR Clerk, since by definition there cannot be any conflict of interest. Otherwise, a containment relationship should not have been used to inherit implicit properties that conflict with explicit properties being defined.

2.4.2 *Role Cardinality.* Another type of constraint imposed on the RBAC authorization database is the cardinality of a role. Some roles in an organization may be occupied by a certain number of employees at any given time. For example, consider the role of a department head. Although over time a number of individuals may assume this role, only one individual may assume the responsibilities of the department head at a given point in time. Cardinality constraints could also be used as a means of enforcing licensing agreements. We formally define cardinality by the following rule: $ML(r : role) \rightarrow \mathscr{X}$, the number of users (0 or more) that may be authorized for role r.

Cardinality. The number of users authorized for a role cannot exceed the role's cardinality:

$$(\forall r : role)\#RM[r] \le ML[r]. \qquad (3)$$

Note that cardinality applies to all authorizations, that is, both explicitly assigned and those resulting from inheritance (see Section 3.1.1).

2.5 Static Properties

The model constraints described in previous sections have a number of important implications for properties of the RBAC model [Kuhn 1997]. The practical significance of the next result is that a role cannot inherit another role that was designated as mutually exclusive to it. This is clearly a
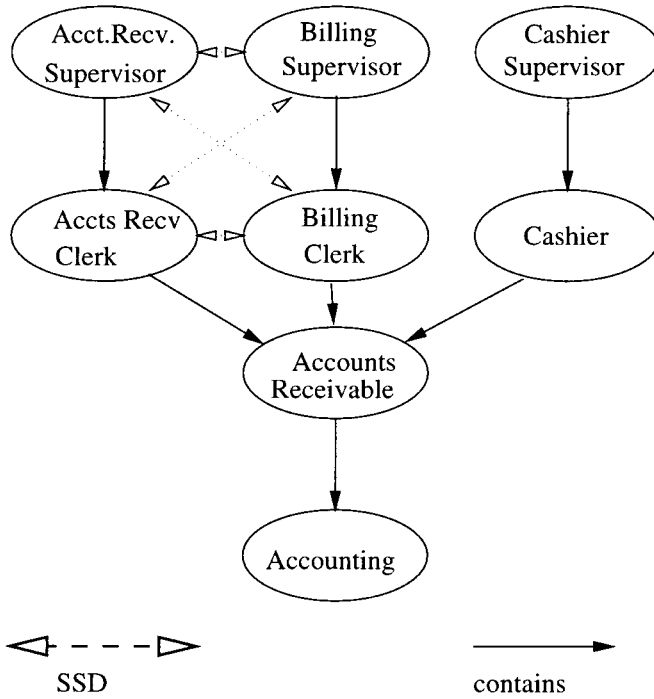
Fig. 4. Example, SSD consistency for AR Clerk is in SSD with Billing Clerk.

desirable property, and this result shows that the rules are sufficient to ensure it.

THEOREM 1. CONSISTENCY OF SSD AND CONTAINMENT. *Two nonempty roles, i and j (i.e., $RM[i] \neq \emptyset$ and $RM[j] \neq \emptyset$) can be mutually exclusive only if they are incomparable within the role hierarchy poset: $(i, j) \in Ea \Rightarrow \neg (i \geq j \vee j \geq i)$.*

PROOF. Suppose $(i, j) \in Ea \wedge (i \geq j \vee j \geq i)$. Arbitrarily choose i as the containing role and let u be authorized for role i. Then, by role inheritance,

$$u \in RM[j].$$

But by SSD, $(i, j) \notin Ea$, which contradicts the assumption. □

An immediate corollary is that if there are any mutually exclusive roles, then a nonempty role cannot be mutually exclusive with itself. This might have been required as one of the basic rules, but as it happens, it is one of their consequences.

Corollary 1. Consistency Among Contained Roles. A nonempty role cannot be mutually exclusive with itself: $\forall i : (i, i) \notin Ea$.

PROOF.   By the theorem above,

$$i \geq j \vee j \geq i \Rightarrow (i, j) \notin Ea.$$

Substituting $j := i$ gives

$$i \geq i \Rightarrow (i, i) \notin Ea.$$

By definition, $i \geq i$, so for all i,

$$(i, i) \notin Ea. \qquad \qquad \square$$

If there are any mutually exclusive roles, then those roles cannot have a common upper bound.

THEOREM 2. NONINHERITANCE OF MUTUALLY EXCLUSIVE ROLES.   *If there is any pair $(i, j) \in Ea$, then there can be no nonempty role k such that $k \geq i \wedge k \geq j$.*

PROOF.   Suppose there is some nonempty role k, and mutually exclusive roles i, j such that

$$k \geq i \wedge k \geq j \wedge (i, j) \in Ea.$$

Then because

$$k \geq i \wedge k \geq j,$$

the role inheritance rule requires that

$$u \in RM[i] \wedge u \in RM[j],$$

so by the SSD rule,

$$(i, j) \notin Ea,$$

which contradicts the assumption. $\square$

An immediate corollary is that the rules also prohibit the existence of a "superuser" or "root" role that contains all other roles on the system.

Corollary 2. Nonexistence of Superuser Role.   If there is any pair $(i, j) \in Ea$, then there can be no nonempty role r such that for all $i, r \geq i$.

Users thus cannot gain access to unauthorized privileges directly through inheritance, but it is also important to show that access cannot be gained indirectly.

THEOREM 3. TRANSITIVITY OF MUTUAL EXCLUSION. *If a user inherits a role that is mutually exclusive with another role, then the user cannot be authorized for the second role. That is,* $u \in RM[i] \wedge i \geq j \wedge (j, k) \in Ea \Rightarrow \notin RM[k]$.

PROOF. Suppose $u \in RM[i] \wedge i \geq j \wedge (j, k) \in Ea \wedge u \in RM[k]$. Then by the Inheritance rule,

$$u \in RM[j].$$

But by SSD,

$$u \in RM[j] \wedge u \in RM[k] \Rightarrow (j, k) \notin Ea,$$

which contradicts the assumption.   □

Mutual exclusion between roles and role cardinality are examples of static constraints, i.e., restrictions on user/role authorization. Inheritance rules apply to arbitrary static constraints $\mathscr{C}(RM[i])$ and $\mathscr{C}(RM[j])$ on users in roles i and j. In contrast to privileges, which can be thought of as inherited "downward," static constraints are, in general, inherited "upward":

THEOREM 4. INHERITANCE OF STATIC CONSTRAINTS. *For roles i, j, if $i \geq j$ and $\mathscr{C}(RM[j])$, where $\mathscr{C}$ is some constraint on users authorized for role j, then $\mathscr{C}(RM[i])$.*

PROOF. If $i \geq j$ and $\mathscr{C}(RM[j])$ then, by role inheritance,

$$u \in RM[i] \Rightarrow u \in RM[j].$$

So by definition of $\in$,

$$RM[i] \subseteq RM[j],$$

so $\mathscr{C}(RM[i])$ holds.   □

As stated earlier, the complexity of dealing with consistency issues is handled by the RBAC software through the implementation of a series of integrity checks that are derived from the RBAC model.

RBAC maintains the integrity of the RBAC authorization database by checking and enforcing the consistency rules described above. For example, administratively imposing Accounts Receivable Clerk in SSD with Billing Clerk, within the hierarchies illustrated in Figure 3, results in the SSD relationships illustrated in Figure 4. That is, through inheritance of static constraints (2.5), the following relationships are guaranteed: AR Supervisor is in SSD with Billing Clerk; Billing Supervisor is in SSD with AR Clerk; and AR Supervisor is in SSD with Billing Supervisor.

2.6 MC_1:RBAC Activation Model

Dynamic properties complement static properties and refer to properties of the model that involve either subjects or mappings from subjects to other basic elements (i.e., subject to user and subject to active-roles). Dynamic properties include role activation, permission execution, dynamic separation of duties, and object access. Dynamic properties provide extended support for the principle of least privilege, in that each user has different levels of permission at different times, depending on the role being performed. These properties ensure that permissions do not persist beyond the time that they are required for performance of duty. This aspect of least privilege is often referred to as *timely revocation of trust*. Revocation of privileges can be a rather complex issue without the facilities of dynamic separation of duty, and has been ignored in the past for reasons of expediency.

In applying dynamic security policy to a computer system, we speak of subjects, which are active entities whose access to roles, operations, and objects must be controlled. All requests by a user are carried out by subjects acting on the user's behalf. Each subject is uniquely referenced by an identifier, which is used to determine whether the subject is authorized for a role and can become active in the role. Role inheritance for the active role set can be handled in a number of ways. The model leaves open the question of whether roles are inherited automatically into the active role set. (See Section 3.3 for the explanation of how RBAC/Web addresses this question.) Subjects are defined by the following type:

  type Subject of subject identifiers

The following functions formalize mappings for users, subjects, and roles:

  $SU(s : subject) \rightarrow user$, the subject to user mapping, which gives the user associated with subject s;

  $AR(s : subject) \rightarrow 2^{role}$, the active role mapping, which gives the set of roles in which subject s is active.

  Role activation. A subject cannot have an active role that is not authorized for its associated user: .

$$(\forall s : subect, u : user, r : roles) r \in AR[s] \Rightarrow SU[s] \in RM[r]. \qquad (4)$$

RBAC is described in terms of a series of mapping functions on the basic types. By considering the subject-user and active-role mapping functions of the model component M_1, along with the mapping functions of the model component M_0, we can now express RBAC in terms of these relationships. These mappings are illustrated in Figure 5 and are used along with a series of constraints on these mappings to express the properties of the model. The solid lines represent static mapping functions among RBAC elements,
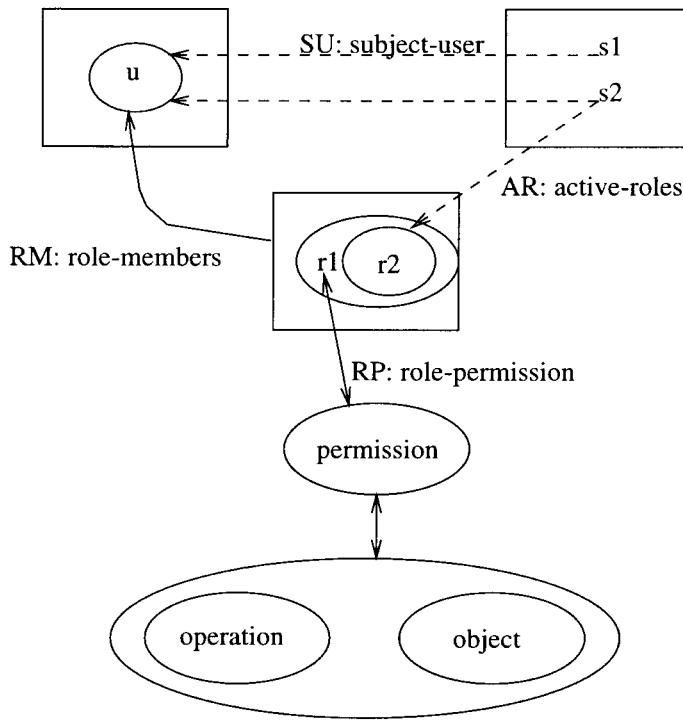
Fig. 5. RBAC functional mappings.

and the dotted lines represent dynamic mapping functions among RBAC elements.

2.7 Dynamic Constraints

With RBAC, administrators can enforce an organization-specific policy of dynamic separation of duty (DSD). With SSD, an organization can address potential conflict-of-interest issues at the time a user's membership is authorized for a role. With DSD, it is permissible for a user to be authorized as a member of a set of roles that do not constitute a conflict of interest when acted on independently, but which produces policy problems when allowed to be acted on simultaneously. For example, a user may be authorized for both the roles of Cashier and Cashier Supervisor, where the supervisor is allowed to acknowledge corrections to a Cashier's open cash drawer. If the individual acting in the role Cashier attempted to switch to the role Cashier Supervisor, RBAC would require the user to drop his or her Cashier role, and force the closure of the cash drawer before assuming the role Cashier Supervisor. As long as the same user is not allowed to assume both roles at the same time, a conflicts of interest situation will not arise. Although this effect could be achieved through an SSD relationship, DSD relationships generally provide the enterprise with greater operational flexibility.

The DSD constraint is formally defined using the following constraint function and property:

$Er : role \times role$ = the set of role pairs (i, j) that are mutually exclusive with each other at activation, that is., no user may be active in both i and j simultaneously

Dynamic separation of duty (DSD). A pair of roles may be designated as mutually exclusive regarding role activation. That is, a user may be active in only one of the two distinct roles so designated: .

$$(\forall s, t : subject)(\forall i, j : role) : s \neq t :$$

$$i \in AR[s] \wedge j \in AR[t] \wedge (i, j) \in Er \Rightarrow SU[s] \neq SU[t]. \tag{5}$$

Note that in the RBAC/Web implementation, dynamic separation of duties is only relative to a user on a single web server (see Section 3.1.2).

2.7.1 *Operation Authorization.* Note that, unlike roles in an SSD relation, roles in a DSD relation can be hierarchically related through the containment relation. This is consistent with the DSD property of restricting simultaneous activation of roles and that of a role hierarchy as a representation of a user's implicit and explicit authorizations for a role. As such, authorization and activation can be treated as independent notions.

In earlier models, authorization and activation were linked, so that if a subject selected a particular role r_i, then any junior roles r_j such that $r_i \geq r_j$ would be brought into the active role set as well. With this approach, if there were any roles r_j with a DSD relationship with r_i, then r_i could not be activated because doing so would require the conflicting roles to be simultaneously active [Kuhn 1997] (see Section 2.5). The model described in this paper does not dictate a specific way to handle role inheritance for active roles. An approach to this problem by means of inheritance and activation hierarchies is described in Sandhu [1998]. An alternative, used in the RBAC/Web implementation for Unix, described in Section 3.3, is to forbid roles that are related hierarchically from having a DSD relationship. This prohibition is enforced when role relationships are added or when DSD relation Er is defined.

$$exec : subject \times operation \rightarrow boolean$$

$$exec(s, op) = \begin{cases} 1 & \text{if subject } s \text{ can execute operation } op \\ 0 & \text{otherwise.} \end{cases}$$

Operation authorization. A subject can execute an operation only if the operation is authorized for the role in which the subject is currently active.

$$(\forall s : subject)(\forall op : operation)$$

$$exec(s, op) \Rightarrow (\exists r : role)(\exists p : permission)r$$

$$\in AR[s] \wedge p \in RP[r] \wedge op \in POp[p]. \tag{6}$$

2.7.2 *Object Access*. To ensure enforcement of enterprise policies for RBAC objects, subject access to RBAC objects must be controlled. The following rule determines if a subject can access an RBAC object:

$$access : subject \times operation \times object \rightarrow boolean$$

$$access(s, op, o) = \begin{cases} 1 & \text{if subject } s \text{ can access object } o \text{ using operation } op \\ 0 & \text{otherwise.} \end{cases}$$

With the role activation property defined above, the object access authorization property defined below ensures that a subject's access to an RBAC object can be achieved through authorized operations by authorized active roles only.

Object access authorization. A subject can perform an operation on an object only if there exists a role that is an element of the subject's active role set and the role contains a permission that authorizes the operation on the object:

$$(\forall s : subject)(\forall o : object)(\forall op : operation)access(s, p, op) \Rightarrow$$

$$(\exists r : role)(\exists p : permission)r$$

$$\in AR[s] \wedge p \in RP[r] \wedge op \in POp[p] \wedge o \in POb[p]. \tag{7}$$

3. RBAC/WEB IMPLEMENTATION

3.1 Design Issues

The NIST RBAC model provides the opportunity for the implementor, based on requirements, to choose between several alternative implementation approaches. Implementation decisions must be made in the design of the administration tools and in the process used to activate roles.

3.1.1 *Administration*. While RBAC can be treated as either a discretionary or nondiscretionary access control method, the treatment in this paper takes the latter approach. That is, one or more administration roles distinct from user roles are required, insofar as their permissions deal solely with the policy attribute elements of the model: user-to-role and role-to-permission mappings, containment relations, cardinality constraints, and separation of duty constraints. Users not assigned to administration roles are denied these permissions and must operate within the confines of the roles defined for them and assigned to them by an administrator.

The administrative model for RBAC/Web follows this approach. There is a single role designated for administration. In the context of the implementation, this role is required to access administrative tools and, consequently, the RBAC database. The administrative model for RBAC/Web is minimal. An example of a richer, more flexible, model can be found in Sandhu et al. [1997].

Another design decision for administration in RBAC/Web is how to manage role authorization when there are role hierarchies. In the NIST RBAC model, if a user is authorized for a role r_0 and r_0 inherits r_1, then that user is also authorized for r_1. What should happen when authorization for r_0 is removed? Should the user retain the authorization for r_1 or should the authorization for r_1 be removed as well?

In RBAC/Web, the design decision was made to remove the authorization for r_1 also. The goal is to make administration in RBAC/Web as easy as possible, while making it difficult for an administrator to make a serious mistake. To accomplish this, the decision was to keep the number of operations in role authorizations to a minimum. Since authorizing a role automatically authorizes the roles it inherits (this is the definition of role inheritance), it is consistent to remove authorizations for inherited roles when authorization for a parent is removed. If the administrator wishes these inherited roles to be authorized for the user, those operations must be done explicitly. With this approach, administrative errors result in roles being unintentionally left unauthorized. The alternative approach, leaving inherited roles authorized, could result in roles being unintentionally left authorized.

Removing authorization for inherited roles when authorization for a parent is removed also results in constantly reminding the administrator of the organization's role hierarchy. To further support this approach, the RBAC/Web Admin Tool differentiates between role *assignment* and role *authorization* as defined in the NIST RBAC model. A role is assigned to a user explicitly by the Admin Tool. A role is authorized as defined in the NIST RBAC model if that role is assigned to the user or is inherited by a role assigned to the user. It follows that if an administrator attempts to assign a role to a user who, by virtue of inheritance from an assigned role, is already authorized for that role, then that attempt is deemed an error.

3.1.2 *Role Activation.* In the NIST RBAC model, there is a function, $AR(s : subject) \rightarrow 2^{role}$, that maps subjects that identify a user (e.g., a login name) to active role sets (ARS) (see Section 2.6). The subject represents a user within the implementation environment. There is also a function in the NIST model, $SU(s : subject) \rightarrow user$, which maps subjects to users; that is, given a subject, the function SU returns the user whom that subject represents. An implementation of the NIST RBAC model must define the concept of *subject* within the context of the implementation environment, the function SU, and the function AR.

Within the environment of a single computer system, e.g., a Unix system, a subject could be defined as the process ID of a user process. The user process represents the user and performs operations in the user's name. In the world wide web environment, each access (e.g., click on a hyperlink) is analogous to the creation of a user process on the web server to perform the access. Once the access completes, the user process is removed.

Due to the short lifetime of user processes within the web server environment, defining the process ID of these processes as the user's subject implies that an active role set must be established upon each access request. Since this is inefficient and unnecessary, RBAC/Web defines the user's subject as the user's login name on the web server. On a web server, each user has only one login name. Thus, there is only one subject per user; that is, the function SU consists of ordered pairs of the form $(loginname, user)$, one for each web server user. Consequently, there is only one active role set per user per web server; that is, the function AR consists of ordered pairs of the form $(loginname, ARS)$, one for each web server user. There is only only one active role set per user per web server, regardless of the location and number of browser windows that may be open to the web server by the user.

RBAC/Web implements this approach by maintaining the user's ARS in a file on the web server. This file has the user's login name, i.e., the user's subject, as part of the file name—thus making it possible to locate a user's ARS, given the user's login name. Each time a user attempts to access a URL, RBAC/Web determines whether the user has access by consulting this file. The RBAC/Web Session Manager, a common gateway interface (CGI) script, part of RBAC/Web, manages the contents of a user's ARS. In order to simplify the process of determining the contents of a user's ARS, when a user has authorized roles that have DSD relationships, the user is presented with a selection of the largest subsets of his or hers set of authorized roles and asked to choose. For the scenario describing how users access a Web server enhanced with RBAC/Web, see Section 3.6.2.

The DSD relationship of the NIST RBAC model places a constraint on active role set contents; that is, any pair of roles in the active role set cannot have a DSD relationship. Given this constraint, the question arises as to how to determine active role set contents when role pairs have both a hierarchical and a DSD relationship.

Consider the following example: suppose role r_1 and role r_2 have a DSD relationship, r_1 is authorized for user U, and r_1 inherits r_2. When establishing U's active role set, the following apparent contradiction results. Role r_2 belongs in U's active role set because r_1 inherits r_2 but r_2 cannot be in U's active role set because r_1 and role r_2 have a DSD relationship. There are at least two possibilities for resolving this in an implementation:

—the inheritance relationship between a pair of roles in DSD is overridden and one role or the other (but not both) is placed in the active role set; or

—the situation, that is, a pair of roles related both hierarchically and having a DSD relationship, is never allowed in the RBAC database.

The design decision made by RBAC/Web for Unix is to implement the latter. In RBAC/Web for Unix, the Admin Tool does not permit a role pair to simultaneously have both a hierarchical and a DSD relationship. Thus, the apparent contradiction in active role set contents can never occur. This design decision is based on the desire for all role relationships specified in the RBAC database to hold at all times and in all situations. The goal is to simplify the task of administration. Administrators are not required to be aware of situation-sensitive rules. They are able to know that the RBAC database holds throughout the administration, role activation, and enforcement of role relationships and access.

Alternative approaches are equally valid, depending on implementation requirements. One such alternative approach is described by Sandhu [1998],which describes the distinction between the *usage* and *activation* aspects of role hierarchies. In terms of these concepts, RBAC/Web's implementation of role hierarchies combines both aspects into one hierarchy model.

3.2 Components

RBAC for the world wide web (RBAC/Web) is an implementation of the NIST RBAC model for web servers on both the internet and intranets. RBAC/Web can be used in conjunction with existing WWW authentication and confidentiality services. These include username/password and Secure Socket Library (SSL). User identification information is passed to RBAC/ Web by the web server. It is the responsibility of the web server to authenticate user identification information and provide confidential data transmission as configured by the web server administrator.

RBAC/Web places no requirements on a browser. Any browser that can be used with a particular web server can be used with that server enhanced with RBAC/Web. RBAC/Web is implemented in two environments: UNIX (e.g., Netscape, Apache servers) and Windows NT (e.g., Internet Information server, website, or purveyor).

3.3 RBAC/Web for Unix

RBAC/Web for Unix implements MC_0 and MC_1, that is, the role properties: hierarchy, cardinality, SSD, and DSD. Components of RBAC/Web are shown in Table I; RBAC/Web for UNIX uses all components in Table I.

With RBAC/Web for UNIX, there are two ways to use RBAC/Web with a UNIX Web server. The simplest is by means of the RBAC/Web CGI. The RBAC/Web CGI can be used with any existing UNIX server without modifying its source code. RBAC URLs are passed through the web server and processed by the RBAC/Web CGI. RBAC/Web configuration files map URLs to file names, while providing access control based on the user roles. Installation of the RBAC/Web CGI is similar to the installation of the web server.

Table I. RBAC/Web Components

Database	Files specifyng relationships between users and roles, role hierarchy, constraints on user/role relationships, current active roles, and relationship between roles and operations.
API Library	Specification used by web servers and CGIs to access RBAC/Web database. The application program interface (API) is the means to add RBAC to any web server implementation. API library is a C and Perl library that implements RBAC/Web API.
CGI	Implements RBAC as a CGI for use with any existing web server without modifying the server. RBAC/Web CGI uses RBAC/Web API.
Session Manager	RBAC/Web Session Manager activates roles by establishing a user's current active role set (ARS).
Admin Tool	Allows server administrators to create users, roles, and permitted operations; associate users with roles and roles with permitted operations; specify constraints on user/role relationships; and maintain the RBAC database. Administrators access the RBAC/Web Admin Tool with a Web browser.

While RBAC/Web CGI is relatively simple to install and use, it is not as efficient as performing access control directly in the web server. The other way to use RBAC/Web is to modify the UNIX Web server to call the RBAC/Web API to determine RBAC access. A URL is configured as an RBAC-controlled URL by means of the web server configuration files that map URLs to file names.

Some web servers for a UNIX environment, such as Netscape and Apache, divide their operation into steps and allow each step to be enhanced or replaced by means of a configuration parameter. This allows web server operations to be modified without having to change the server's source code. For these web servers, the RBAC/Web API can be integrated by simply providing the appropriate calling sequence and modifying configuration parameters.

3.4 RBAC/Web for Windows NT

RBAC/Web for Windows NT only implements MC_0, that is, the static properties: hierarchy, cardinality, and SSD. In order to implement DSD in a Windows NT environment, it is necessary to change the session establishment mechanisms of Windows NT. Such a difficult task was beyond the scope of the project. The task is particularly difficult in that neither documentation nor source code for Windows NT session establishment is public information. Because RBAC/Web for Windows NT only implements

MC_0 and because built-in NT security mechanisms are closely compatible with RBAC, RBAC/Web for Windows NT needs only the database and administrative tool components shown in Table I. RBAC/Web for Windows NT requires no modification of Web server internals or access to source code.

Because Windows NT security mechanisms are uniform between Windows NT itself and many Web servers that run on Windows NT, RBAC/Web for Windows NT can also manage access to the Windows NT File System as well as Web resources. In addition to the Admin Tool which manages user/role and role/role relationships, the tool, RGP-Admin [Barkley and Cincotta 1998], was implemented to manage role/permission relationships in the Windows NT environment.

3.5 RBAC Database

The RBAC database is an implementation of the RBAC model component MC_1 for RBAC/Web for Unix and MC_0 for RBAC/Web for Windows NT. The RBAC database elements are created using the administrative tool that graphically displays and maintains these relationships. In addition, the RBAC administrative tool maintains the integrity of the RBAC database by checking and enforcing consistency rules[Gavrila and Barkley 1998].

An operation as defined in Section 2.2 represents an access method to a set of one or more protected RBAC objects. When authorizing user membership into a role, the user is implicitly provided with the potential for exercising operations associated with the role. Permissions in the RBAC/Web are HTTP methods that an end-user can perform on RBAC-controlled URLs.

In general, constraints provide confidence as to the adherence of enterprise-wide policies. In theory, similar effects can be achieved by establishing procedures and sedulous actions of administrators. For example, administrators can maintain and share a list of role pairs known to be mutually exclusive and ensure that an individual user is never authorized for both roles. However, the reality is that procedures break down and administrators get reassigned over time. The constraints imposed by RBAC/Web provide management and regulators with the confidence that critical security policies are uniformly and consistently enforced within the network and, as such, contribute to the network's operational assurance.

Associated with objects managed by RBAC/Web are the ACLs that reside with each web server. With RBAC/Web for Unix, an ACL is organized as a list of roles, where, for each role, there is a list of HTTP methods under which a user acting in the role is permitted to access an associated URL. RBAC/Web for NT makes use of the ACLs built into Windows NT where a role maps to a Windows NT domain or local group. The collection of ACLs is organized and managed as the collection of role-privilege relationships.

In the context of RBAC/Web, each subject represents a user active in one or possibly many roles. A user establishes an active role set, that is, a subset of the roles in which the user has membership is activated. A user's

Fig. 6. RBAC/Web Admin Tool: Role hierarchy display.

authorization (a consequence of role membership) is a necessary, but not always sufficient, condition for a user to be permitted to execute a privilege. Other organizational policy considerations or constraints that pertain to authorizing users to execute permissions may need to be taken into consideration.

RBAC/Web requires that a user be authorized as active in a role before being permitted to perform an operation or access a URL. This provides the context for imposition of other policy checks. In the case of RBAC/Web for Unix, administrators can enforce an organization-specific dynamic separation of duty (DSD) policy.

3.6 Use Scenario

There are two use scenarios. One describes how an administrator manages access policy; the other describes end-user interaction.
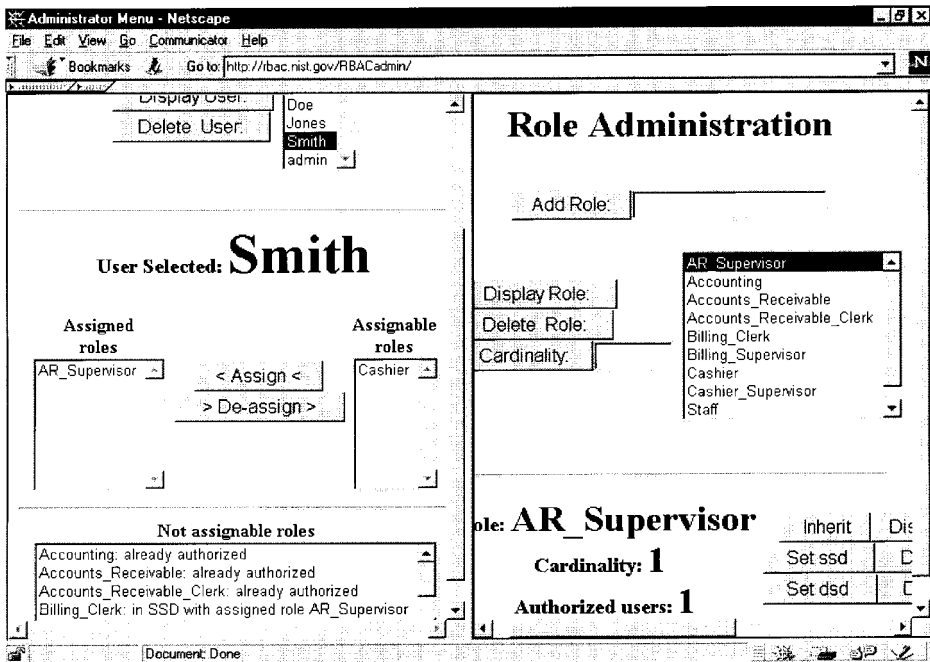
Fig. 7. RBAC/Web Admin Tool: Display for managing user/role and role/role relationships.
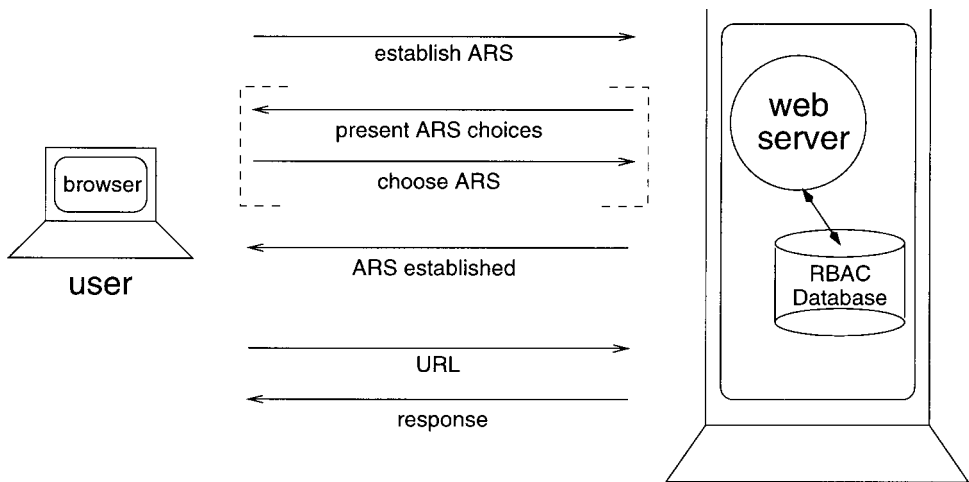

Fig. 8. RBAC/Web end-user perspective.

3.6.1 *Use by Administrators*.   The RBAC/Web Admin Tool manages us-
er/role and role/role associations. The RBAC/Web Admin Tool manages the
specification of user/role assignments, role hierarchies, static separation of
duty constraints, dynamic separation of duty constraints (in the case of
RBAC/Web for Unix), and role cardinality.

In order to reduce errors in administration, Admin Tool differentiates between role *assignment* and role *authorization* as defined in the NIST RBAC model. A role is assigned to a user explicitly by Admin Tool. A role is authorized as defined in the NIST RBAC model if that role is assigned to the user or is inherited by a role assigned to the user. Role assignment helps an administrator maintain awareness of the role hierarchies that describe an organization.

The Admin Tool manages a RBAC database by means of a fixed set of operations that can be performed on the database. A formal description of these operations can be found in Gavrila and Barkley [1998].

RBAC/Web is used to manage access to the Admin Tool itself. There is a single special administrative role that must be authorized for a user to access the Admin Tool and RBAC database.

The example access control policy for an accounting department, shown in Figure 4, is illustrated in the Admin Tool screen captures of Figures 6 and 7. The role hierarchy display in Figure 6 presents a graphical view of the sample accounting department policy. In particular, it shows the accounting department's role hierarchy. Figure 6 and the role relationships of the selected role AR Supervisor (shaded) show that the role AR Supervisor has a SSD relationship with the roles Billing Supervisor and Billing Clerk. These SSD relationships are indicated by enclosing the roles Billing Supervisor and Billing Clerk in pentagons instead of ovals. The parenthesized expressions, (n, m), under each role name in Figure 6, indicates the cardinality of each role (m) and how many users the role was authorized for (n). The "U" indicates unlimited cardinality, i.e., there is no limit on the number of users who may be authorized for the role.

The left frame of Figure 7 shows the display for managing user/role assignments. The frame shows that user Smith was assigned the role AR Supervisor. The frame also shows that, in concert with the role/role associations defined by accounting department policy as shown in Figure 6, Smith could also be assigned the role Cashier. However, Smith could not be assigned the roles Accounting, Accounts Receivable, or Accounts Receivable Clerk because these roles are inherited by the role AR Supervisor; nor could Smith be assigned the role Billing Clerk or Billing Supervisor because these roles have an SSD relationship with the role AR Supervisor. These role/role associations can be seen using the graphical display shown in Figure 6.

The right frame of Figure 7 shows the display for managing role/role associations. It presents buttons for establishing and removing all role/role associations. The frame shows that the role AR Supervisor was selected. Since the role AR Supervisor was selected in the right frame of Figure 7, the Role Hierarchy display of Figure 6 also shows the role AR Supervisor selected.

Admin Tool enforces a set of consistency requirements on the RBAC/Web database. These consistency requirements, described in Gavrila and Barkley [1998], ensure that user/role and role/role relationships have the

required relationships to each other. For example, as described in Section 3, RBAC/Web for Unix requires that any roles that have a DSD relationship cannot have a hierarchical relationship.

3.6.2 *Use by End-Users*.   From the end-user's perspective, because DSD is not supported, interaction with a web server enhanced with RBAC/Web for Windows NT is exactly the same as interacting with any other Windows NT web server. However, with RBAC/Web for Unix, before access to a URL controlled by RBAC is permitted, end-users must establish an active role set (ARS) as shown in Figure 8. End-users choose and/or are assigned a current active role set. The ARS determines the HTTP methods that the end-user can perform on RBAC-controlled URLs. The ARS remains in effect until the end-user establishes a new ARS. It is the RBAC/Web session manager that allows users to change their ARSs.

A user may be assigned roles that have DSD relationships. If this is the case, the session manager enables users to choose the subset of their assigned role set that they would like to have in their ARS. Users are presented with a list of subsets that do not violate any DSD relationships and asked to choose. In order to minimize the number of choices, the subsets in the list, taken from the set of all possible subsets of a user's assigned roles, contains the largest subsets that do not violate any DSD relationships. Once the choice is made, an ARS consisting of all assigned roles in the chosen subset and all roles that the assigned roles inherit is established. If there are no DSD relationships among the roles assigned to a user, then the ARS consisting of all authorized roles is automatically established. Note that because the RBAC/Web implementation applies only to a single server, dynamic separation of duties is not implemented relative to users across multiple servers (e.g., in an environment where a collection of servers constitutes an administrative security domain), but only relative to a user on a single web server (see Section 3 for how RBAC/Web implements a user's ARS on a web server).

4.   CONCLUSIONS

Although intranets can offer great benefits to a company or government agency, security problems remain. For intranets to reach their full potential in enterprise computing, access control mechanisms must be in place that can conveniently, and cost effectively, regulate user access to information, while providing management with confidence that their critical policies are faithfully and consistently enforced across administrative boundaries. To solve these and other authorization problems, NIST has initiated an effort to provide and promote the use of role-based access control (RBAC) for intranet web servers. RBAC is particularly attractive for intranet applications because it can reduce the complexity and cost of authorization management. In addition, RBAC provides a context for the specification and enforcement of complex security policies that are often impractical or even impossible to enforce through the direct use of conventional access control mechanisms.

RBAC is administered through roles and role hierarchies that mirror an enterprise's job positions and organizational structure. Users are assigned membership into roles consistent with a user's duties, competency, and responsibility. To address conflicts of interest issues, constraints are imposed on user membership in roles and on a user's ability to activate a role. Complexities introduced by simultaneously supporting mutually exclusive roles and role hierarchies are handled by the RBAC software, making security administration easier. Roles, role hierarchies, and constraints provide the context with which the intranet administrators can specify, and RBAC/Web servers can enforce, the specifics of a large variety of laws, regulations, and business practices.

RBAC supports several well-known security principles and policies important to commercial and government enterprises that process unclassified but sensitive information [Ferraiolo et al. 1993; van Solms and van der Merve 1994]. These include specification of competency to perform particular tasks; the enforcement of least privilege for administrators and general users; and the specification, as well as enforcement, of conflicts of interest rules, which may entail duty assignment and dynamic and static separation of duties. For RBAC/Web, these policies can be enforced at the time that users are authorized as members of a role, at the time of role activation (e.g., when a role is established as part of a user's active session), or at the time the user attempts to access a URL.

Under RBAC, intranet administrators are provided with a single view of the RBAC authorization database at a level of abstraction that is intuitive and consistent with the way the enterprise is structured and conducts business. RBAC/Web thereby bridges the huge gap between enterprise laws, regulations, and business practices and the details of the underlying access control mechanisms of web servers.

REFERENCES

BARKLEY, J. AND CINCOTTA, A. 1998. Managing role/permission relationships using object access types. In *Proceedings of the 3rd ACM Workshop on Role-Based Access Control* (RBAC, Fairfax, VA, Oct. 22-23). ACM Press, New York, NY, 73–80.

FERRAIOLO, D. AND KUHN, D. R. 1992. Role based access control. In *Proceedings of the 15th Annual Conference on National Computer Security*. National Institute of Standards and Technology, Gaithersburg, MD, 554–563.

FERRAIOLO, D., CUGINI, J., AND KUHN, D. R. 1995. Role based access control: Features and motivations. In *Proceedings of the 11th Annual Conference on Computer Security Applications*. IEEE Computer Society Press, Los Alamitos, CA, 241–248.

FERRAIOLO, D. F., GILBERT, D. M., AND LYNCH, N. 1993. An examination of federal and commercial access control policy needs. In *Proceedings of the 16th National Conference on Computer Security* (Baltimore, MD, Sept. 20–23). National Institute of Standards and Technology, Gaithersburg, MD, 107–116.

FEINSTEIN, H. L. 1995. Final report: NIST small business innovative research (SBIR) grant: Role based access control: Phase 1. SETA Corporation. SETA Corporation.

GAVRILA, S. AND BARKLEY, J. 1998. Formal specification for role-based access control user/role and role/role relationship management. In *Proceedings of the 3rd ACM Workshop on Role-Based Access Control* (RBAC, Fairfax, VA, Oct. 22-23). ACM Press, New York, NY, 81–90.

HOFFMAN, J. 1997. Implementing RBAC on type enforced systems. In *Proceedings of the 13th Annual Conference on Computer Security Applications*. IEEE Computer Society Press, Los Alamitos, CA, 158–163.

JANSEN, W. A. 1988. Revised model for role based access control. NIST-IR 6192. National Institute of Standards and Technology, Gaithersburg, MD.

KUHN, D. R. 1997. Mutual exclusion as a means of implementing separation of duty requirements in role-based access control systems. In *Proceedings of the 2nd ACM Workshop on Role-Based Access Control* (Fairfax, VA, Nov. 6-7). ACM Press, New York, NY, 23–30.

NYANCHAMA, M. AND OSBORN, S. L. 1994. Access rights administration in role-based security systems. In *Proceedings of the IFIP Working Group 11.3 Working Conference on Database Security*. Elsevier North-Holland, Inc., Amsterdam, The Netherlands.

SANDHU, R., COYNE, E. J., FEINSTEIN, H. L., AND YOUMAN, C. E. 1996. Role-based access control models. *IEEE Comput. 29*, 2 (Feb.), 38–47.

SANDHU, R. AND MUNAWER, Q. 1998. How to do discretionary access control using rules. In *Proceedings of the 3rd ACM Workshop on Role-Based Access Control* (RBAC, Fairfax, VA, Oct. 22-23). ACM Press, New York, NY, 47–54.

SANDHU, R. 1998. Role activation hierarchies. In *Proceedings of the 3rd ACM Workshop on Role-Based Access Control* (RBAC, Fairfax, VA, Oct. 22-23). ACM Press, New York, NY, 33–42.

SANDHU, R., BHAMIDIPATI, V., COYNE, E., GANTA, S., AND YOUMAN, C. 1997. The ARBAC97 model for role-based administration of roles: Preliminary description and model. In *Proceedings of the 2nd ACM Workshop on Role-Based Access Control* (Fairfax, VA, Nov. 6-7). ACM Press, New York, NY, 41–54.

SIMON, R. AND ZURKO, M. E. 1997. Separation of duty in role based access control environments. In *Proceedings of the 10th IEEE Workshop on Computer Security Foundations* (Rockport, MA, June 10-12). IEEE Computer Society Press, Los Alamitos, CA, 183–194.

VON SOLMS, S. H. AND VAN DER MERVE, I. 1994. The management of computer security profiles using a role-oriented approach. *Comput. Secur. 13*, 8, 673–680.