



Role-based access control in DCOM

Gail-Joon Ahn

Department of Computer Science, University of North Carolina at Charlotte, NC 28223, USA

Received 4 June 1999; received in revised form 14 February 2000; accepted 27 May 2000

Abstract

The explosive growth of the Web, the increasing popularity of PCs and the advances in high-speed network access have brought distributed computing into the mainstream. To simplify network programming and to realize component-based software architecture, distributed object models have emerged as standards. One of those models is distributed component object model (DCOM) which is a protocol that enables software components to communicate directly over a network in a reliable, and efficient manner. In this paper, we investigate an aspect of DCOM concerning software architecture and security mechanism. Also, we describe the concept of role-based access control (RBAC) which began with multi-user and multi-application on-line systems pioneered in the 1970s. And we investigate how we can enforce the role-based access control as a security provider within DCOM, specially in access security policy. © 2000 Elsevier Science B.V. All rights reserved.

1. Introduction

Distributed applications introduce new design and deployment issues. For this added complexity to be worthwhile, there has to be a significant payback. Some applications are inherently distributed. Multiuser games and teleconferencing applications are examples of such applications. For these, the benefits of a robust infrastructure for distributed computing are obvious. Many other applications are also distributed, in the sense that they have at least two components running on different machines. But because these applications were not designed to be distributed, they are limited in scalability and ease of deployment. Any kind of workflow or groupware application, most client/server applications, and even some desktop productivity applications essentially control the

way their users communicate and cooperate. Adopting the notion of distribution into these applications benefits the user and optimizes the use of network and computer resources. The application designed with distribution in mind can accommodate different clients with different capabilities by running components on the client side when possible and running them on the server side when necessary.

DCOM is an extension of the component object model (COM). COM [6] defines how components and their clients interact. This interaction is defined such that the client and the component can connect without the need of any intermediary system component. The client calls methods in the component without any overhead whatsoever. This study will look at the DCOM architecture and DCOM's security mechanism. DCOM provides an extremely efficient default security mechanism that lets developers write distributed applications without having to worry about security

E-mail address: gahn@uncc.edu (G.-J. Ahn).

at all. Any security provider supported by windows NT can be used with DCOM's security mechanism.

The concept of role-based access control (RBAC) began with multi-user and multi-application on-line systems pioneered in the 1970s. The central notion of RBAC is that permissions are associated with roles, and users are assigned to appropriate roles. This simplifies management of permissions. Roles are created for the various job functions in an organization and users are assigned roles based on their responsibilities and qualifications. Users can be easily reassigned from one role to another. Role can be granted new permissions as new applications and systems are incorporated, and permissions can be revoked from roles as needed. Several researchers have showed that RBAC can be accommodated in current systems, such as Oracle, Unix. To enforce role-based access control in DCOM with advantages that RBAC allows us to achieve, we use DCOM's group mechanism which is similar to the notion of roles in RBAC. Even though DCOM does not have role concept inside, it is possible for RBAC to be applied to DCOM security mechanism, which can satisfy DCOM's access security policy and achieve ease of administration.

This paper begins with the description of the DCOM model in Section 2. In Section 3, we discuss the RBAC model followed by Section 4 which includes integration between RBAC model and DCOM. Section 5 concludes the paper.

2. Distributed component object model

In programming and engineering disciplines, a component is an identifiable part of a larger program or construction. Usually, a component provides a particular function or group of related functions. In programming design, a system is divided into components that in turn are made up of modules. In object-oriented programming and distributed object technology, a component is a reusable program building block that can be combined with other components in the same or other computers in a distributed network to form an application. Examples of a component

include: a single button in a graphical user interface, a small interest calculator, an interface to a database manager. Components can be deployed on different servers in a network and communicate with each other for needed services [10,3]. A component runs within a context called a container. Examples of containers include pages on a Web site, Web browsers, and word processors.

COM is Microsoft's framework for developing and supporting program component objects. The COM provides a set of interfaces allowing clients and servers to communicate within the same computer (running a Windows 95 or NT system). It is aimed at providing similar capabilities to those defined in Common Object Request Broker Architecture (CORBA), the framework for the interoperation of distributed objects in a network. Whereas OLE provides services for the compound document that users see on their display, COM provides the underlying services of interface negotiation, life cycle management (determining when an object can be removed from a system), licensing, and event services (putting one object into service as the result of an event that has happened to another object).

DCOM is a protocol that enables software components to communicate directly over a network in a reliable, and efficient manner. At the same time it is a program interface in which client program objects can request services from server program objects on other computers in a network.

For example, one can create a page for a Web site that contains a script or program that can be processed (before being sent to a requesting user) not on the Web site server but on another, more specialized server in the network. Using DCOM interfaces, the Web server site program (now acting as a client object) can forward a remote procedure call (RPC) to the specialized server object, which provides the necessary processing and returns the result to the Web server site. This result is passed on to the Web page viewer.

In following section, we will look at the DCOM's architecture and the processing steps DCOM takes. We will also discuss its security mechanism and policy.

2.1. DCOM's architecture

A client that needs to communicate with a component in another process cannot call the component directly, but has to use some form of inter-process communication provided by the operating system. COM provides this communication in a completely transparent fashion: it intercepts calls from the client and forwards them to the component in another process. When client and component reside on different machines, DCOM simply replaces the local inter-process communication with a network protocol. Neither the client nor the component is aware that the wire that connects them has just become a little longer. Fig. 1 shows the overall DCOM architecture: the COM run-time provides object-oriented services to clients and components and uses RPC and the security provider to generate standard network packets that conform to the DCOM wire-protocol standard.

Table 1 describes how DCOM works [8]. In summary, the client program calls `CoInitialize`, then `CoCreateInstance` passing the CLSID and IID(s). The client receives back Interface Pointer(s), and then calls the desired method. The server program implements the interface as a class, and has only two extra calls to `CoInitialize` and `CoRegisterClassObject`.

2.2. DCOM's security mechanism

Different platforms use different security providers, and many platforms even support multiple security providers for different usage scenarios or for interoperability with other platforms [1,12]. DCOM and RPC are designed in such a way that they can simultaneously accommodate multiple security providers. All these security providers provide a means of identifying a security principal, a means of authenticating a security principal, and a central authority that manages security principals and their keys. A security principal is a module that is associated with access privileges in the systems. It represents an entity such as individual user or a group. Every time a user logs in to the system it is as a particular principal. Each user may have several principals associated with the user. On the other hand each principal is required to be associated with a single user. This list of security principals is called an access control list (ACL). Access control lists consist of multiple access control entries (ACE), which correspond to individual security principals. Each access control entry can indicate that the corresponding security principal is to be allowed access or that it is to be denied access. If a client wants to access a secured resource, it passes its security identity and some form of authenticating data to the resource and the resource asks the security provider to authenticate the client. Security providers typically use low-level custom protocols to interact with clients and protected resources.

Using the network for distributing an application is challenging not only because of the physical limitations. It also raises new issues related to security between and among clients and components. Since many operations are now physically accessible by anyone with access to the network, access to these operations has to be restricted at a higher level. Without security support from the distributed development platform, each application would be forced to implement its own security mechanisms. A typical mechanism would involve passing some kind of username and password – usually encrypted – to some kind of logon method. The application would validate these credentials against a user database or directory and return

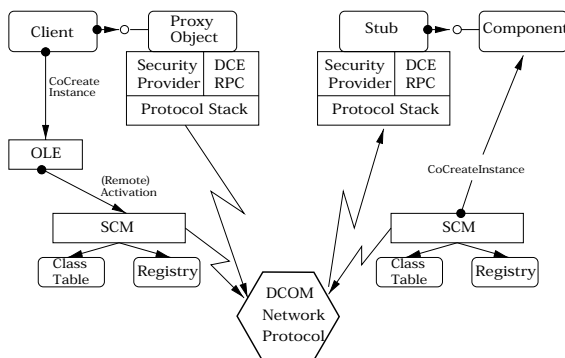


Fig. 1. DCOM architecture.

Table 1
DCOM steps

Client initialization

The client calls CoInitialize which initializes the COM library for use. Next, the client calls COM library's CoCreateInstance, passing the ID for the desired class (CLSID) and an array of desired interface IDs (IIDs).

Server activation (Client side)

The call to CoCreateInstance is passed on to the Service Control Manager (SCM) on the client machine which looks up the Class ID in the local Class Table to see if the server is already running. If not the SCM will look in the Registry using the CLSID to find the type and location of the server.

Server activation (Server side)

Using the server machine address, the client SCM establishes an RPC link with the server machine's SCM, passing it the desired CLSID and IID(s). The server SCM checks its Class Table to see if the server is running, and if it is not, it looks up the CLSID in the Registry to find out what command needs to be executed to start the server, and executes it. When a COM server starts up, the first thing it calls is CoInitialize. The second thing it calls is CoRegistryClassObject, passing the implemented CLSID and a pointer to the server program's Class Factory. This effectively advertises the server program to DCOM, adding its entry into the Class Table.

The Class Factory

Whenever a client requests a new instance of a server class, the SCM calls a method called CreateInstance on the server program's Class Factory. This is the common gateway used by all clients. CreateInstance then instantiates the object and passes the pointer to the object's IUnknown interface back to the SCM.

Multiple query interface (MQI)

Once the SCM has a pointer to the object's IUnknown interface, it can use QueryInterface to request pointers to other interfaces.

Proxy/Stub Loading

Data passed back between client and server have to be packaged into RPC packets. This process is called marshalling, and is performed by the Proxy/Stub DLL(s). Using the Interface ID (IID) returned from the MQI step, the SCMs on each machine interrogate the Registry for the DLL(s) that need to be loaded into the client's and server's process spaces. Interface Pointers are then handed back to the client that issued the CoCreateInstance call.

Method Call

Using the Interface Pointer, the client calls the server method.

some dynamic identifier for use in future method class. On each subsequent call to a secure method, the clients would have to pass this security identifier. Each application would have to store and manage a list of usernames and passwords, protect the user directory against unauthorized access, and manage changes to passwords, as well as dealing with the security hazard of sending passwords over the network. A distributed platform must thus provide a security framework to safely distinguish different clients or different groups of clients so that the system or the application has a way of knowing who is trying to perform an operation on a component.

DCOM uses the extensible security framework provided by Windows NT. Windows NT provides a solid set of built-in security providers that support multiple identification and authentication mechanisms, from traditional trusted-domain security models [5] to noncentrally managed, massively

scaling public-key security mechanisms [2]. A central part of the security framework is a user directory, which stores the necessary information to validate a user's credentials. DCOM can make distributed applications secure without any security-specific coding or design in either the client or the component. Just as the DCOM programming model hides a component's location, it also hides the security requirements of a component. The same binary code that works in a single-machine environment, where security may be of no concern, can be used in a distributed environment in a secure fashion [7].

DCOM's default security mechanism is illustrated in Fig. 2. DCOM achieves security transparency by letting developers and administrators configure the security settings for each component. Just as the Windows NT File System lets administrators set access control lists (ACLs) for files and directories, DCOM stores access control lists

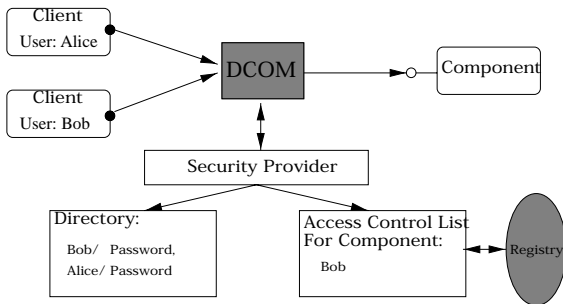


Fig. 2. Security mechanism.

for components. These lists simply indicate which users or groups of users have the right to access a component of a certain class. These lists can easily be configured using the DCOM configuration tool. Whenever a client calls a method or creates an instance of a component, DCOM obtains the client's current username associated with the current process. Windows NT guarantees that this user credential is authentic. DCOM then passes the username to the machine or process where the component is running. DCOM on the component's machine then validates the username again using authentication mechanism and checks the access control list for the component. If the client's username is not included in this list (either directly or indirectly as a member of a groups of users), DCOM simply rejects the call before the component is ever involved. This default security mechanism is completely transparent to both the client and the component and is highly optimized. It is based on the Windows NT security framework, which is probably one of the most heavily used parts of the Windows NT operating system: on each and every access to a file or even to a thread-synchronization primitive like an event or semaphore.

2.3. DCOM's security policy

DCOM distinguishes four fundamental aspects of security [4]:

- *Access security: protecting the object.* Which security principals are allowed to call an object? The most obvious security requirement on dis-

tributed applications is the need to protect objects against unauthorized access. Sometimes only authorized users are supposed to be able to connect to an object. In other cases, non-authenticated or unauthorized users might be allowed to connect to an object, but must be limited to certain areas of functionality. Current implementations of DCOM provide declarative access control on a per-process level.

- *Launch security: protecting the server machine.* Which security principals are allowed to create a new object in a new process? Another related requirement on a distributed infrastructure is to maintain control over who can create objects. Since all COM objects of a machine are potentially accessible via DCOM, it is critical to prevent unauthorized users from creating instances of these objects.
- *Security identity: controlling the object.* What is the security principal of the object itself? Another aspect of distributed security is that of controlling the objects themselves. Since an object performs operations on behalf of arbitrary callers, it is often necessary to limit the capabilities of the object itself. One obvious approach is that of making the object assume the identity of the caller. Whatever action the object performs is limited by the caller's privileges. Although managing access can be simplified by using user groups, it is often simpler to have the object itself run under a dedicated security identity, independent of the security identity of the current caller.
- *Connection policy: integrity and privacy.* DCOM gives both callers and objects a range of choices to determine how the data on the connection are to be secured. DCOM offers two fundamental choices with regard to data protection: integrity and privacy. Clients or objects can request that data be transferred with additional information that ensures data integrity. If any portion of the data is altered on its way between the client and the object, DCOM will detect this and automatically reject the call. Data integrity implies that each and every data packet contains authentication information.

The scope of this paper will be within the first security policy, *access security*. The following

sections will show how to simulate this security issue with RBAC.

3. Overview of RBAC model

RBAC is an alternative policy to traditional mandatory access control (MAC) and discretionary access control (DAC) [20,21]. As MAC is used in the classical defense arena, the policy of access is based on the classification of objects such as top-secret level [15]. The main idea of DAC is that the owner of an object has discretionary authority over who else can access that object [9,19]. But RBAC policy is based on the role of the subjects and can specify security policy in a way that maps to an organization’s structure.

A general family of RBAC models called RBAC96 was defined by Sandhu et al. [18]. Fig. 3 illustrates the most general model in this family. Motivation and discussion about various design decisions made in developing this family of models are given in [16,18]. Also, there are variations regarding distributed systems [22].

Fig. 3 shows (regular) roles and permissions that regulate access to data and resources. Intuitively, a user is a human being or an autonomous agent, a role is a job function or job title within the organization with some associated semantics regarding the authority and responsibility conferred on a member of the role, and a permission is an approval of a particular mode of access to one or more objects in the system or some privilege to

carry out specified actions. Roles are organized in a partial order \geq , so that if $x \geq y$, then role x inherits the permissions of role y . Members of x are also implicitly members of y . In such cases, we say x is senior to y . Each session relates one user to possibly many roles. The idea is that a user establishes a session and activates some subset of roles that he or she is a member of (directly or indirectly by means of the role hierarchy). The RBAC model has the following components and these components are formalized from the above discussions.

- U is a set of users,
- R is disjoint sets of roles and administrative roles, respectively,
- P is disjoint sets of permissions and administrative permissions,
- $UA \subseteq U \times R$, is a many-to-many user to role assignment relation,
- $PA \subseteq P \times R$ is many-to-many permission to role assignment relations,
- $RH \subseteq R \times R$ is partially ordered role hierarchies (written as \geq in infix notation),
- S is a set of sessions,
- $user : S \rightarrow U$, is a function mapping each session s_i to the single user $user(s_i)$ and is constant for the session’s lifetime,
- $roles : S \rightarrow 2^R$ is a function mapping each session s_i to a set of roles $roles(s_i) \subseteq \{r \mid (\exists r' \geq r)[(user(s_i), r') \in UA]\}$ (which can change with time) so that session s_i has the permissions $\bigcup_{r \in roles(s_i)} \{p \mid (\exists r'' \leq r)[(p, r'') \in PA]\}$, and
- there is a collection of constraints stipulating which values of various components of the RBAC model are allowed or forbidden.

A user can be a member of many roles and a role can have many users. Similarly, a role can have many permissions and the same permissions can be assigned to many roles. Each session relates one user to possibly many roles. Intuitively, a user establishes a session during which the user activates some subset of roles that he or she is a member of the permissions available to the users are the union of permissions from all roles activated in that session. Each session is associated with a single user. This association remains constant for the life of a session. A user may have multiple sessions open at the same time, each in a

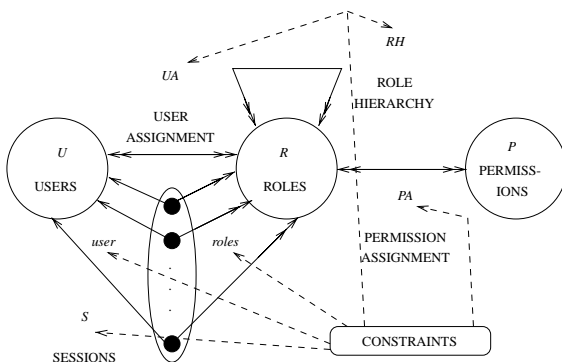


Fig. 3. RBAC model.

different window on the workstation screen for instance. Each session may have a different combination of active roles. The concept of a session equates to the traditional notation of a subject in access control. A subject is a unit of access control, and a user may have multiple subjects (or sessions) with different permissions active at the same time.

4. Integration with RBAC

In this section we outline one approach to enforcing RBAC in DCOM. For some applications, a single component-wide access control list which is described in Section 2.2 is not sufficient. Some methods in a component may be accessible only to certain users. For example, an accounting business component may have a method for registering new transactions and another method for retrieving existing transactions. Only members (such as Bob) of the accounting department (user group “Accounting”) should be able to add new transactions, while only members (such as Alice) of transaction management (user group “Transaction”) should be able to view the transactions. And members of

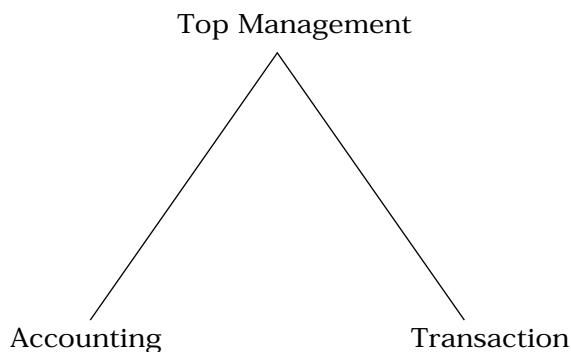


Fig. 6. An example of role hierarchy.

top management (user group “Top Management”) should be able to add and view the transaction. How can an application use DCOM security to implement the selective security required in this example?

We can approach this example with programmatic control using DCOM [11]. This approach is shown in Fig. 4. When a method call comes in, the component asks DCOM to impersonate the client. After this, the called thread can perform only those operations on secured objects, that the client is permitted to perform. The component can then try to access a secured object, such as a registry key, that has an ACL on it. If this access fails, the client was not contained in the ACL, and the component rejects the method call. By choosing different registry keys according to the method that is being called, the component can provide selective security in a simple way.

We can simulate RBAC in DCOM along the lines of this example¹. Fig. 5 shows a conceptual abstraction of a role-based access control model. With this abstraction we can see that RBAC can be inserted into DCOM architecture (in Fig. 1) as a part of security provider.

In order to use RBAC96, we should accommodate role hierarchies. Roles would map to NT user group(s) (which do not support hierarchies). For example, user group “Accounting” would map to role “Accounting”. We can represent the

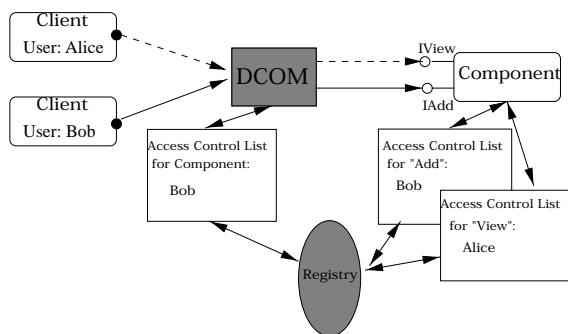


Fig. 4. Security using Registry Key.




Fig. 5. Conceptual abstraction of RBAC.

¹ We assume that the assignment is done by Sandhu and Ahn [13,14], including the simulation of role hierarchies.

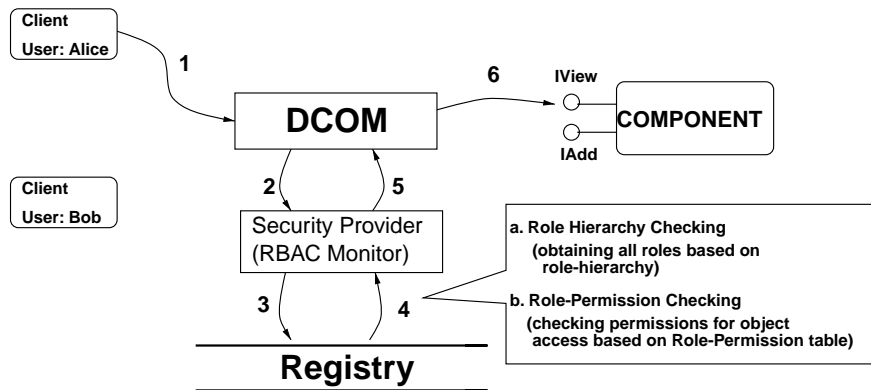


Fig. 7. Integration with RBAC.

role-relationship as role hierarchy. The role hierarchy of this example would be as shown in Fig. 6. The *Accounting* role can have permission to add new transactions, while only the *Transaction* role can have permission to view the transactions. The *Top Management* role is senior to *Accounting* and *Transaction* and thereby inherits all permissions from junior roles.

In Fig. 4, the DCOM checks the ACL of component and then the component checks the ACL of methods. Whenever the component is accessed the ACL of the component should be checked. Instead of doing this two-step process, we can have a unified checking mechanism using RBAC as shown in Fig. 7. For ease of reference, we call the portion of security provider that enforces role-based access control as RBAC monitor. At the role hierarchy checking step, the client's all roles that he or she is a member of (by means of role hierarchies) should be checked. After this step RBAC monitor decides whether the client can have permission(s) to access the component according to given role(s) during the role-permission checking.² For example, assume that a client Chris has a role *Top Management* in Fig. 6 and tries to access an accounting business component. During role hierarchy check, we can know Chris's role membership: {*Accounting*, *Transaction*, *Top*

Management}. After the role permission check, he can have all permissions from junior roles such as *Accounting* and *Transaction* using Registry Key. This means Chris can add new transactions and also view the transaction. Let us consider that Bob has a role *Accounting* in Fig. 6 and tries to access an accounting business component. We can simply know that Bob can only add new transactions but he cannot view the transaction. Using an approach with RBAC we can have the same result as the previous approach illustrated in Fig. 4. The RBAC approach can also reduce the interactions between a component and secure object such as Registry.

In order to enforce our framework described in this paper, we also need to construct role-permission relation. For example, consider the medical system which has role-permission table illustrated in Table 2, and deals with critical information such as medical records. Prescription, Test-Result, Medical Record, and Financial Record are objects which include the information accessed by users who have permissions to access those objects. And each permission is an interface for each object. Table 2 implies that *Physician* role can have *read* and *write* permissions to objects such as Prescription and Medical Record, and he/she can just read the object Test-Result. And *Patient* role can just read only his/her objects such as Prescription, Medical Record, and Test-Result. That is, this table displays the relation between roles and permissions. Even though this table is not the real example, it is sufficient to show the need of role-

² These permissions can easily be configured using DCOM configuration tool (addressed in Section 2.2).

Table 2
An example of role-permission table

	Prescription	Test-Result	Medical Record	Financial Record
Physician	R W	R	R W	–
Nurse	–	–	R	–
Pharmacist	R	–	–	–
Accountant	–	–	–	R W
Patient	R	R	R	–

permission table. This kind of access policy allows us to prevent an unauthorized user from doing malicious thing on (with) the critical information.

Integrating with RBAC, we just showed that we can simplify access control in DCOM. Using our approach, we can enforce role-based access control that can authorize each access to objects in DCOM. In further study, we will investigate the details how RBAC can be achieved in the end-to-end mechanism of DCOM.

5. Conclusion

In this paper, we have described the architecture and security mechanism of DCOM. Also, we briefly looked at the RBAC model as a security specification model. Finally, we have shown that RBAC can be accommodated in DCOM. Also we can see that the adoption of RBAC as a part of security provider can simplify access control and provide administrative convenience. This framework was just based on DCOM's access security policy. In the future work, we also would investigate whether DCOM has sufficient flexibility to accommodate administrative access control models such as URA97 [17].

References

- [1] P.E. Ammann, J.C. Knight, Data diversity: an approach to software fault-tolerance, *IEEE Transactions on Computers* 37 (4) (1988) 418–425.
- [2] C. Adams, S. Lloyd, *Understanding Public-key Infrastructure: Concepts, Standards, and Deployment Considerations*, Macmillan, New York, November 1999.
- [3] K.P. Birman, *Building Secure and Reliable Network Applications*, Manning Publications, 1996.
- [4] Microsoft Corporation, DCOM architecture, in: *Proceedings of Microsoft Professional Developers Conference*, September 1997.
- [5] Department of Defense National Computer Security Center, *Department of Defense Trusted Computer Systems Evaluation Criteria*, DoD 5200.28-STD, December 1985.
- [6] D. Rogerson, *Inside COM*, Microsoft Press, 1996.
- [7] J. Gray, A. Reuter, *Transaction Processing: Concepts and Techniques*, Morgan Kaufman, Los Altos, CA, 1993.
- [8] R. Grimes, *Professional DCOM Programming*, Wrox Press, 1997.
- [9] S. Jajodia, P. Samarati, V.S. Subrahmanian, E. Bertino, A unified framework for enforcing multiple access control policies, in: *Proceedings of the ACM SIGMOD International Conference on Management of Data Pages*, 1997, pp. 474–485.
- [10] M. Moriconi, et al., *Secure software architectures*, in: *Proceedings of IEEE Symposium on Research in Security and Privacy*, IEEE Press, New York, 1997.
- [11] Microsoft Corporation, *Microsoft Windows NT Server: DCOM Technical Overview 1995*, www.microsoft.com/com/dcom95.
- [12] B. Randell, System structure for software fault tolerance, *IEEE Transactions on Software Engineering SE-1* (12) (1995) 220–232.
- [13] R. Sandhu, G.-J. Ahn, Decentralized group hierarchies in UNIX: an experiment and lessons learned, in: *Proceedings of 21st NIST-NCSC National Information Security Conference*, Crystal City, VA, October 1998, pp. 486–502.
- [14] R. Sandhu, G.-J. Ahn, Group hierarchies with decentralized user assignment in Windows NT, in: *Proceedings of IASTED Conference on Software Engineering*, Las Vegas, NV, October 1998, pp. 352–355.
- [15] R.S. Sandhu, Lattice-based access control models, *IEEE Computer* 26 (11) (1993) 9–19.
- [16] R. Sandhu, Rationale for the RBAC96 family of access control models, in: *Proceedings of the First ACM Workshop on Role-Based Access Control*, ACM, New York, 1997.
- [17] R. Sandhu, V. Bhamidipati, The URA97 model for role-based administration of user-role assignment, in: T.Y. Lin, X. Qian (Eds.), *Database Security XI: Status and Prospects*, North-Holland, Amsterdam, 1997.
- [18] R.S. Sandhu, E.J. Coyne, H.L. Feinstein, C.E. Youman, Role-based access control models, *IEEE Computer* 29 (2) (1996) 38–47.

- [19] R. Sandhu, Q. Munawar, How to do discretionary access control using roles, in: *Proceedings of the Third ACM Workshop on Role-Based Access Control*, 1998, pp. 47–54.
- [20] R.S. Sandhu, P. Samarati, Access control: principles and practice, *IEEE Communications* 32 (9) (1994) 40–48.
- [21] R. Sandhu, P. Samarati, Authentication, access control and intrusion detection, in: E. Bertino (Ed.), *The Computer Science and Engineering Handbook*, CRC press, Boca Raton, 1997, pp. 1929–1948.
- [22] N. Yiaelis, E. Lupu, M. Sloman, Role-based security for distributed object systems, in: *Proceedings of the IEEE Fifth Workshops on Enabling Technology: Infrastructure for Collaborative Enterprise*, IEEE, New York, 1996.



Gail-Joon Ahn is an assistant professor of Computer Science Department at University of North Carolina, Charlotte, North Carolina. His principal research and teaching interests are in information and systems security. He received his Ph.D. and M.S. degrees from George Mason University, Fairfax, Virginia, and B.S. degree in Computer Science from SoongSil University, Seoul, Korea. He was a research associate at the Laboratory for Information Security Technology, George Mason University. His research interests include access control, security architecture for distributed objects, and secure e-commerce systems. He is a member of ACM and IEEE Computer Society.