A simple RBAC implementation

GSEC

Practical Assignment (1.4c)

Version 5

Hotze de Jong
5 October 2004

Table of Contents

Abstract

In this document a practical case is described of an organization in which access to confidential information is governed in an automated way. The organization produces and develops technical complex products.

Confidential information consisting of trade secret information and customer confidential information must be shared by the employees in several departments. There is a need-to-share information opposed to a need-to-know.
A system was required to manage the access rights to this kind of information. The organization was not ready for the implementation of a full blown RBAC-based system, so another more simple system was required.

In this document a system is described based on an access matrix which governs the access to confidential information for groups of people.

Confidential information is defined based on the type of document. Groups are defined as creating groups and reading groups based on common attributes owned by the HR department. A matrix is introduced that gives reading groups access to confidential information of creating groups.
An escape procedure is given for people who need access that is not provided to the group they are in. These people are added to a guest lists that provides access to the needed confidential information.

The document starts with a short description of the situation before the access matrix was implemented. Next the problem is described.
Then the four aspects of the solution are discussed: classification of information, grouping people, the access matrix itself and the escape method.

The current situation describes a number of experiences with this system as it is used for some years now.

A comparison to the core RBAC model is made, resulting in the conclusion that this system provides the functionality as described for the core RBAC model.

Finally the conclusion gives some advantages and disadvantages of the implemented system.

Introduction

Within the company the development information was stored in several directories, both personal and project/department directories. Information stored in these directories was secured based on the group concept in Unix; sometimes properly restricted but most were open to the world.

It was possible to publish html pages with confidential content and links to documents on the company intranet. No authorization structure was available on the intranet so access to confidential information could not be restricted. All users had by default access to the intranet.

Risk management dictated a more restrictive way to access confidential information.

A project to develop a document management system was started. The specifications stated that a way to restrict access to confidential documents should be part of the system. This required a method to manage the access to confidential information. This resulted in the initiation of the access matrix project.

This document starts with a problem analysis. Four separate problems were defined that are described below: classifying information, grouping people, granting access and providing in an escape method.

Problem description

In the company are two different sets of information:

1. Information that is managed in specific applications like SAP and Oracle applications. These applications have their own authorization model. Furthermore different kinds of information are contained in these systems ranging from part numbers and module prices to planning lots. Segregation of duty and the principle of least privilege are important aspects in these systems because they are used to manage the business processes in the company.
2. Information contained in development documents, specifications and product related documents, etc. This information is product related and not process related. An important aspect of this information is that groups of people need to have access. Information needs to be readily available otherwise some developer will invent the same trick once again (reinvent the wheel), thus wasting time and money.

The method described in this paper discloses the second set of information.

Because this information, at least part of it, is classified as trade secret or customer confidential, the access must be restricted to those people who need to have access. In fact this kind of information can be described as need-to-share information opposed to need-to-know information. Need-to-know is more restrictive, there must be a valid reason to get access, whereas the need-to-share method is more concerned with ensuring that all people have access to the information they may need.

The problem is described as: How can we provide access to confidential information in an easy way and still restrict access to those who need to have access?

Two important aspects are:

1. Because access to information must be managed, segregation of duties is no requirement. Segregation of duties [1] is required in processes were care must be taken that a single individual has no control over two or more phases of a transaction or operation.
2. Access to confidential information must be provided in a secure way. However, because the information must be shared it is no requirement that only access to information that is needed is provided (need-to-know). The result is that the principle of least privilege is not applicable, although the privileges must be kept as reduced as possible.

Requirements

To solve this problem the next requirements were stated.

1. The information must be classified in an easy way. If possible based on an algorithm that can be implemented in an automated system.
2. The groups that need to have access to confidential information must be managed in an easy way, preferable in an automated way.
3. The access to confidential information must be described in a simple model that can be managed in an easy way.
4. An escape method must be provided if the above model is too restrictive for certain individuals.

The next conditions were defined:

- The information must not be over-classified. Over classification may slow down the business process due to the extra precautions required for secure handling and storage. Information that is over-classified will soon cause employees to disregard the classification system, rendering organization information protection programs ineffective.
- The classification of information must not be subject to personal interpretation, but must be clearly defined on agreed rules.
- All information must have an information owner who determines who may have access to her or his confidential information.
- Authorization of access to confidential information must not be based on persons but on functions or other groups of employees. Individual authorizations cannot be maintained in an effective and efficient way. The number of different groups must be kept to a minimum.
- To reduce the risks concerned with confidential information, access to confidential information must be based on a business need. Groups that have do not need access to the information must be excluded.
- Authorization of access to confidential information must preferably not be done on the content of the information but on the type of information. Classification on content can only be done by individuals and will lead to personal interpretation of classification rules.
- The classification system must be designed in such a way that it is possible to simply implement and maintain it in an automated system.

Classification of information

All information within the company is classified according to the classification policy [2], [3].
There are two categories of information:

1. Unclassified information
   All information that is available for the public like the published financial report of a company.
2. Classified information
   Classified information is information that must be kept within the organization and to which access is restricted to a particular class of people. The desired degree of secrecy about such information is known as its sensitivity.
   Within classified information three classes can be distinguished based on the sensitivity:
   - Secret information; this is very sensitive information, only a very limited group of people will have access. Access is given to individuals only. Special arrangements must be made for this kind of information.
   - Confidential information. This group contains trade secret information (sensitive in the sense that the competitiveness of the company is dependent on it), customer confidential information (information that can be related to specific customers) and other confidential information like HR information. Access to confidential information must be based on business need and must be restricted.
   - Information for internal use only; this kind of information is available for everyone within the company.

All information must be classified based on the above four classes.
However:

- Secret information is only created in specific parts of the organization were they know how to deal with it.
- Unclassified information can not be created by normal employees. Only the communications department is allowed to create this kind of information.

Access to these two classes of information is already covered.
Priority must be given in making a distinction between confidential information and information for internal use only (both more than 80% of the total information within the company).

Within the organization a formal development and production process is defined. This process ensures that the development and production of products is well organized and that it produces the required documentation and deliverables. The process consists of several steps. At each step in the process information and documents are generated. The information and documents produced in the first part of the process will generally be more sensitive than that of the last part of the process.

A development process can for instance start with a marketing study, then produce high level specifications, detailed specifications, test plans, etc. The marketing study and high level specifications will be confidential, whereas the rest of the process will produce less sensitive information.
Documents like specifications, studies, test plans, test results are distinguished based on type of document. The type of document indicates the process step in which it is created and thus indicates the sensitivity of the information.
All documents are checked-in in the document management system. The creator must provide the type of document at check-in. Based on the type of document the classification of the document can be determined by the system.

Sometimes specific types of documents are created in more than one step of the process. In this case both the type of document and the step in the process determine the classification of the document.

The next table describes for three departments (marketing, development and production) which type of document they create and how this document is classified.

Type of document	Marketing	Development	Production
Study	Confidential	Confidential	For Internal use only
Specification	Confidential	Confidential	-
Test report	-	For Internal use only	For Internal use only
Acceptance report	-	Confidential	For Internal use only
Business plan	Confidential	-	-
Product roadmap	For Internal use only	-	-
Service contract	-	-	For Internal use only

The marketing department creates four kinds of documents: studies, specifications, business plans and product roadmaps. Three are classified as

confidential and only the product roadmap is classified as ‘for internal usage only’. Test reports, acceptance reports and service contracts are not created by the marketing department.
In this example there must be three information owners; within marketing, development and production. They must determine what the classification of the information must be and who will get access to their confidential information.

Of course this method can only be applied on documents that are checked in in the document management system. Other information, like HTML pages, does not have an associating information type. This kind of information is classified by the person who creates it. This poses a risk because people are subjective in classification or even tend to forget to classify information. Guidelines are provided to support this classification process. Regular audits are performed to check the classification of this kind of information.

Grouping people

As stated in the requirements, the groups that need to have access to confidential information must be managed in an easy way, preferable in an automated way.

In this case only two activities had to be distinguished in the authorization schema: read access and create/delete/modify access. Consequently two kinds of groups can be distinguished: reading groups and creating groups. The creating group is always the owner of the information it creates. The reading group just uses that information.

Grouping people based on individuals is not a way to solve this problem. These kinds of groups tend to grow, because there is no need for the owner to remove people from the list.
In the ideal situation people will be added to or removed from groups in an automated way, reflecting the organization. To do so requires attributes that are attached to all people within the organization. Furthermore their value must be of very good quality.

Groups are created based on HR attributes. These are available for all people within the organization and in this organization HR data is of high quality.
The department of an employee is used as one of the attributes to group on. In this way an organizational reassignment (effected through a change in department) can be reflected in the access to confidential information because this can result in membership of another creating or reading group. To enable this each department gets two extra attributes: the creating group to which it belongs and the reading group to which it belongs. Based on these two attributes every person in a department is appointed to a creating and reading group.
Sometimes it is needed to give managers special access. In this case the position in the organization is used. Special rules are implemented to realize this. Future developments may use the function level also.
In order to make this group concept not to complicated, a person can be member of only one creating and reading group.

In this way membership of a group is based on rules and is dynamically changed based on the value of the attributes.

The next table is an example of grouping people.

Group	Criteria
Marketing creating group	Dept 123, 234 and 345
Marketing reading group	Dept 123, 234 and 345

Group	Criteria
Development creating group	Dept 111, 222, 333, 444 and 555
Development reading group 1	Dept 111 and 222
Development reading group 2	Dept 333, 444 and 555
Production creating group	Dept 678, 789
Production reading group management	Managers of dept 678, 789
Production reading group rest	Dept 678, 789

People in a creating group are able to check-in documents in the document management system. Based on the creating group and the document type the classification is determined by the system.

Defining access: the access matrix

The concept of an access matrix is also described in [2], [4] en [5]. About the same concept is used in this approach. However, the groups of objects and subjects differ in this document.

Two sets of groups have been defined: groups of confidential information and groups of people. The groups of people come in two flavors: creating groups and reading groups. The creating group creates confidential information and is responsible for defining who has access to their information. Creating groups must be created in such a way that all information of a creating group is alike and may be accessed by the same groups of people. The model does not provide a solution for creating groups that create two classes of information with different reading groups; in this case the creating group must be split in two.

An important aspect of the information and groups of people was their homogeneity. For the information this is because the information is grouped on type, but also because it all deals with the product. For the groups of people this is because the departments are created in such a way that they form a homogeneous group that performs about the same activities.
This homogeneity makes it possible to provide access to several types of information at once to relatively large groups of people.

Now a matrix can be drawn in which the access of reading groups to the information of creating groups is defined.

	Creating groups		
Reading groups	Marketing	Development	Production
Marketing	x	x	-
Development group 1	x	x	-
Development group 2	-	x	x
Production management	x	x	x
Production rest	-	-	x

The above matrix shows that people in the marketing department have read access to confidential information of their own department and also to that of the development department. They don't have access to confidential information of the production department. Of course a creating group has always create/delete/modify access to their own information.

To setup the access matrix, a discussion must be organized between the information owners of the different creating groups. To simplify the discussions,

each information owner also represents the needs of the reading groups in his/her department. In the discussion the requirements of the reading group must be matched with the conflicting requirements of the creating group. The information owner of the creating group will try to keep the information as confidential as possible, and give only access to people who need it. On the other hand the information owner of the reading group will ask as much access it can get.
In this way each cell in the matrix must be discussed and a mutual agreement must be reached.

On the intranet people can publish information in two ways: as confidential information or for internal use only. The confidential information is published in a special way with their department as creator. Access to confidential information that is published in this way is also governed by the access matrix.
'For internal use only' information is available for every employee who can access the intranet.

In this way a rather easy way is established to govern the access of groups of people to the confidential information of other groups when the information needs to be shared. This is done by keeping the aggregation level (and so the homogeneity) of the groups as high as possible.

Escape: the guest list

The access matrix governs access for groups of people. There will always be people in these groups who do not have sufficient access to the confidential information they need to fulfill their job.

If a discussion is started about not having sufficient access, two methods can be used to solve this issue.

First of all the access matrix can be adjusted in such a way that the access is provided. But if only a few people require extra access, it is not good practice to give the total group access. So this method should only be used if the major part of the group requires extra access.

If only a few individual people need extra access, the principle of the guest list can be used. Each reading group can have a guest list attached. The people on the guest list get the same access rights as the people in the reading group the guest list is attached to, so by choosing the right guest list individual problems can be solved.

As each guest list can give access to the information of several information owners, in principle all the information owners need to agree if a person is added to a guest list. In practice the decision to put a person on a guest list is delegated to one information owner who decides for all the other information owners.

The disadvantage of the guest list is that membership is not based on rules and will not dynamically change based on the value of the attributes. Membership is fixed; people must be removed from the guest list in order to remove access rights.

To be sure that the guest list does not grow too large, audits must be performed on a regular basis. Automatic clean-up of the guest list is enforced by putting an end date to the membership of each member.

Current situation

The process of implementation of the access matrix was finished some years ago. The access matrix is still working to full satisfaction.
Problems were (and still are)

- Coping with organizational changes. These changes occur quite frequently, the organization is constantly adjusted to the demands of the environment and the management. New departments must get a reading group and a creating group appointed. This cannot be enforced because the access matrix is not managed in the same system as the organization structure. This is solved in two ways:
  1. Regular lists are made of all departments together with the reading and creating group. All departments without these attributes can easily be checked and corrected. This list is made twice a year.
  2. People will complain when they don't have access to confidential information. The first check is always on the attributes of the department. Most departments are corrected in this way.
- A requirement for the access matrix to work optimal is that the organizational structure of the company reflects the need for access to confidential information. If departments are created in such a way that they consist both of people that don't need access to confidential information and a group that does, a potential problem is created. This can be solved in three ways:
  1. have the department split in two (very hard to realize..)
  2. give extra access to the group that doesn't need the access
  3. put the group of people that needs extra access on a guest list.

  Most of the time option 3 must be chosen.

The guest list is audited three times per year. The size is quite constant as are the persons on the list. Occasionally some persons are added. It is very rare that people are deleted from the guest list, even at audit time. The only time when people are deleted from the list is when their end date expires and the person does not complain.

A problem not directly related to the access matrix is the classification of information that is published on the intranet. This must be done by the creator. Because the information is in free format like HTML pages and GCI scripts, there is no system that can perform the classification in an automated way. Information that is classified as confidential can be published in such a way that the access matrix rules the access. But the information is not always properly classified. So regular audits are performed on published information and if needed corrective actions are performed.

The amount of effort to maintain the access matrix is minimal. Changes on the guest list are done by the helpdesk. Changing attributes of departments is done by a special group of people, but occurs only a few times per year; at audit time and with reorganizations.

Comparing the access matrix with the core RBAC model

In this chapter a comparison is made between the implementation of the access matrix and the RBAC model as described in [6], [7] and [8].

Basic data elements

The core RBAC model includes five basic data elements called users, roles, objects, operations, and permissions. The next table describes the way these basic data elements are implemented in the access matrix.

	RBAC	Access matrix
Objects	This can be any system resource subject to access control, such as a file, printer, terminal, database record, etc.	These are the documents and HTML pages to which access must be provided.
Operations	An executable image of a program, which upon invocation executes some function for the user.	There are two operations: reading and creating.
Permissions	An approval to perform an operation on one or more RBAC protected objects	There are two permissions: - permission to create objects - permission to read objects.
Role	A job function within the context of an organization with some associated semantics regarding the authority and responsibility conferred on the user assigned to the role.	There are three roles: - membership of department X. - management of department X. - guest to readinggroup X These roles are translated into reading groups and creating groups that have the relation with the permissions.
User	A human being.	Users are people that are registered in the HR system.

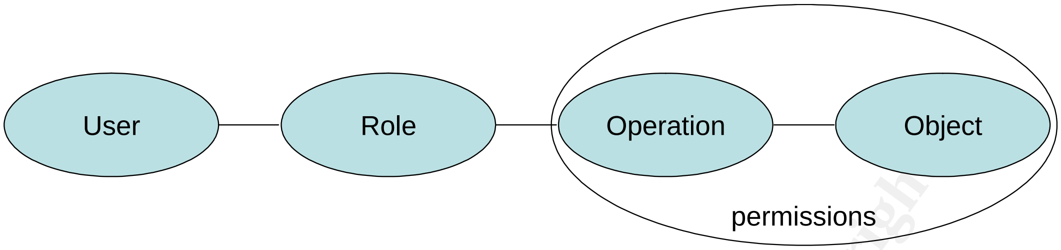

Core RBAC Object model

The core RBAC objectmodel is the same model that is used for the implementation of the access matrix.

The role model is implemented in the next way.

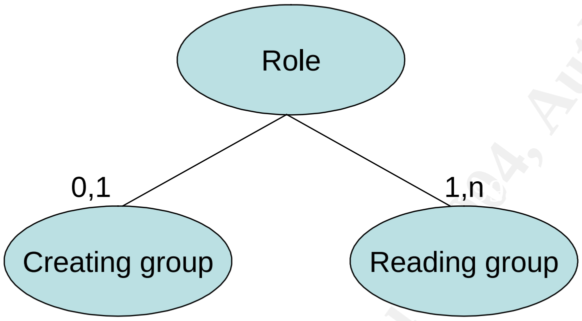

Role model as implemented in the access matrix

Basic functions

Three groups of functions can be found in an RBAC system, see [8] pp 7-9, 15-19:

1. administrative functions
2. supporting system functions
3. review functions

The implementation of the administrative functions within the access matrix is given below.

There are a number of basic administrative commands for a core RBAC model. These can also be recognized in the access matrix. The implementation for the reading and creating groups is for a large part automated. Guest lists are separate lists within the access matrix, but in fact they are just additions to reading groups. There are no RBAC commands specific for the guest list.

Command	Access matrix
Add user, delete user	Based on the HR system. The delete user action is triggered by the end of contract date.
Add role, delete role	Adding or deleting a role for a reading or creating group must be done by hand.
Assign User, De-assign User	Membership of a role is automated based on the HR attributes department and management position. Making a person member of a guest list results in the assignment of the user to this role. This is a manual action. De-assignment is also done manually or triggered by the end date.
Grant Permission, Revoke Permission	The permission for a role to perform an operation on an object are arranged by scripts or ACL's. Changes (new permissions and revoked permissions) have to be done by hand in changing the script or ACL.

For the core RBAC model there are also four supporting system functions defined: Create Session, Add Active Role, Drop Active Role, Check Access. These functions are not implemented in the access matrix. The access control mechanism of the document management system and the ACL's in Intranet take care of these functions.

The review functions within an RBAC system are implemented in a separate audit procedure. Reports are defined to generate the needed information. These functions are:

- Assigned Users
- Assigend Roles

As can be seen from the above comparison, the access matrix performs the same functions as a core RBAC implementation according to the NIST definition, so the access matrix is in fact a simple implementation of the core RBAC model. Its simplicity is based on the homogeneity of the groups of people and the information represented by an information type and on the need-to-share principle.

Conclusion

Main conclusion is the the access matrix was implemented successfully and that it performs the basic functions of a core RBAC model.

To conclude this discussion a comparison is made between a number of implementation aspects access matrix and an RBAC implementation.

The advantages of the access matrix with respect to an RBAC based systems are:

- Scalability
  The groups are managed based on common attributes. Number of employees is no issue. Neither is the amount of information that must be accessed. Experience has learned that organizational changes and growth has been processed in the access matrix without problems. The size of the access matrix itself could become a problem, because if this grows too large the overview gets lost. In reality this has not yet become a problem.
- Cost.
  The costs concerned with the implementation of the access matrix are very limited compared with the implementation of an RBAC based system.
- Ease of implementation.
  The implementation of the access matrix was rather easy. Some scripts were made and the document management system and intranet were adjusted to cope with the new access control method. SAP-HR was adjusted in order to base access rights on HR attributes.
- Transparency and ease of use.
  Only three concepts are used: rules to define confidential information, user groups and an access matrix. This results in a transparent implementation. Users know the access matrix and most of the time they know why they don't have access to confidential information. Maintenance is very simple due to the limited number of aspects to manage.
  Only the guest list makes it somewhat less clear because it is hidden for normal users.
- Flexibility to deviate from the model.
  The guest list provides an easy way to deviate from the groups that are defined.

Disadvantages are:

- Possibility of to much access to confidential Information
  Because groups of people are defined based on common attributes, it is possible that some people have too much access. Compared with the risk this provides and the ease of maintenance and low cost, this risk has been accepted.

When an RBAC based system is implemented the access to confidential information can of course easily be managed by the RBAC system. However, the access matrix has proven to be an easy way to control access to confidential information and provides in this respect a good alternative for a traditional RBAC-based system.

References

1. “Segregation of Duties”. The university of Utah
URL: http://www.utah.edu/Internal_Audit/segregation_of_duties.htm

2. Kurzban, Stanley, “Handbook of Information Security Management“. Chapter 1-3-1 Implementation of Access Controls
URL: http://www.cccure.org/Documents/HISM/081-085.html#Heading6

3. “Security in the Government Sector”, Chapter 3: Information classification, Department of the prime minister and cabinet, 2002
URL: http://www.security.govt.nz/sigs/html/chapter3.html

4. Feng, Hwei-Hsin. Le, Anh. Scanlon, Jeff. “Access matrix”
URL: http://cne.gmu.edu/itcore/security/policy.html

5. “Information access matrix”
URL: http://www.bristol.ac.uk/ISC/bs7799/infomatrix.pdf

6. “An introduction to Role Based Access Control”. NIST/ITL Bulletin, December, 1995
URL: http://csrc.nist.gov/rbac/NIST-ITL-RBAC-bulletin.html

7. Weber, Hazen A. “Role-Based Access Control: The NIST Solution”. GSEC, 8 October 2003
URL: http://www.sans.org/rr/papers/56/1270.pdf

8. “Role Based Access Control”, Secretariat Information Technology Industry Council (ITI). DRAFT - 4/4/2003,
URL: http://csrc.nist.gov/rbac/rbac-std-ncits.pdf