# SECURING HEALTHGRID ENVIRONMENTS

Christos Ilioudis, Dimitrios Baltatzis, George Pangalos
*Informatics Laboratory, Computers Division, Faculty of Technology, Aristotle University of Thessaloniki*
*54006, Thessaloniki, Greece*

Christos Georgiadis
*Department of Applied Informatics, University of Macedonia, 156 Egnatia Street*
*GR-540 06, Thessaloniki, Greece*

Keywords: Authorizations, HealthGrid environment, Mobile agent systems, RBAC.

Abstract: Grid technologies promise to change the way that health organizations tackle complex problems by offering unprecedented opportunities for resource sharing and collaboration. Healthgrids are Grid infrastructures comprising applications, services or middleware components that deal with the specific problems arising in the processing of biomedical data. Resources in Healthgrids are databases, computing power, medical expertise and even medical devices. Securing this new environment in Health organizations is a major issue today. Security considerations and more specifically authorization decisions is a critical problem. Personal data is confidential, so access to the information must be restricted to authorized and authenticated persons. Furthermore data must be protected to guarantee its confidentiality and integrity. This work provides a suitable authorization mechanism that facilitates the usage of grid and agent technology in HealthGrid environments. More specifically, our approach applies the RBAC access control model for dynamically assigning security roles to visiting agents on hosts of the HealthGrid environment. Our methodology proposes a flexible role decomposition method, which facilitates the role assignment process. The role decomposition relies on a set of common Attribute Fields, shared between Grid's hosts, filled with Attribute values that every host evaluates according to its security goals. In any case, every host participating in the grid retains its security policy without altering or compromising its security policy in order to participate in the agent exchange process. The proposed process and the related assignment algorithms have been experimentally implemented and applied in a typical health environment. The results have shown that the proposed framework is applicable and implementable, and can be applied successfully in real life health care environments.

## 1 INTRODUCTION

Healthgrids represent an innovative use of emerging information technology to support broad access to rapid, cost-effective and high quality healthcare. Grid computing aims at the provision of a global Information and Communication Technology infrastructure that will enable a coordinated, flexible, and secure sharing of diverse resources, including computers, applications, data, storage, networks, and scientific instruments across dynamic and geographically dispersed organizations and communities (Virtual Organizations). Goals of HealthGrid infrastructures can be to provide researchers with a broad spectrum of data for analysis and to make it possible for patients to gain access to their data regardless where they are and where it is stored.

HealthGrid systems evolved as an important new field, distinguished from conventional distributed health systems in the sense that they involve large-scale resource sharing guided by innovative applications. Grids provide the necessary infrastructure by which participant organizations across widely distributed systems form and operate for resource sharing (Scott et all, 2004). The infrastructure provides an efficient utilization of existing resources by combining heterogeneous and widely distributed computational resources into a consistent environment. The central idea is that



Ilioudis C., Baltatzis D., Pangalos G. and Georgiadis C. (2007).
SECURING HEALTHGRID ENVIRONMENTS.
In *Proceedings of the Second International Conference on Security and Cryptography*, pages 394-401
DOI: 10.5220/0002123203940401

computing and information sharing should be reliable, pervasive and transparent for widely distributed systems. Securing such Grid infrastructures requires therefore suitable security models and mechanisms.

In contrast, agent technology focuses on the development of concepts, methodologies and algorithms for autonomous problem solving engaged in uncertain and dynamic environments in order to achieve their objectives (Jennings 2001). Agents are actually software entities that can move code, data and state to remote hosts. An agent has the ability to migrate from one host another in order to fulfil its task. Their late ability to move computations across the nodes of a widely distributed network makes them an attractive paradigm, compared to the traditional client-server paradigm. It can be said however that the two technologies attend to service a common environment that is, communities of large distributed systems bound together by a common goal or cause. They also face overlapping problems as Grids seek to become more flexible and agile while agent systems seek to become more reliable and scalable. Various researchers have already observed this. As noted for example by Foster et al "For Grids to be effective in their goals, they must be imbued with flexible, decentralized decision making capabilities while agents need a robust distributed computing platform that allows them to discover, acquire, federate and manage the capabilities necessary to execute their decisions" (Foster et al, 2004).

It is clear that in medical data processing privacy protection needs to be enforced much more rigorously than in other application areas. The users of such systems (patients, medical staff, etc.) are not usually trained in computer security and thus need easy to use and mostly transparent services. Securing HealthGrid infrastructures therefore requires suitable security models and mechanisms with those characteristics, in order to support authentication, data integrity, private communications and access control.

This paper discusses the overall security problem of HealthGrids and focuses on an authorization policy for HealthGrid participant organizations, which utilizes the advantages offered by both the grid and the agent systems technology. We propose a dynamic role assignment mechanism for Grid participants and more specifically we describe an authorization mechanism, which facilitates the use of shared resources in such dynamic environments. Every participant organization in the Grid preserves and maintains its own local security policy, while users active in the Grid retain the ability to access common resources by acquiring local roles for authorization. The proposed approach takes into account and exploits the dynamic characteristics of Grid systems and the flexibility of Role Based Access Control Policies (Ferraiolo et al, 2001). In our approach Grid participants that own resources can specify the authorization policy using a well defined access control language like the eXtensible Access Control Markup Language (XACML) (OASIS 2003) and the Security Assertion Markup Language (SAML 2003). Grid users can also specify their identity and security constraints in the same manner.

# 2 THE SECURITY PROBLEM OF HEALTHGRID INFORMATION SYSTEMS

As already noted, security is an essential consideration when accessing the shared resources of a HealthGrid system. The main security requirements for Grid systems that influence the definition of the HealthGrid security requirements are related to the following characteristics of a Grid (Welch et al, 2003), (Simpson et al, 2006).

## 2.1 Access Control Requirements

The different resources in a grid may have different access policies, including how they authenticate and authorize users. However, if there are no common or overlapping authorizations among the resources, they do not form a usable Grid. Grid service requests can span multiple security domains. Trust relationships among these domains play an important role in the outcome of such end-to-end traversals. A service needs to make its access requirements available to interested client entities, so that they understand how to securely request access to it. Trust between end points can be presumed, based on topological assumptions, or explicit, specified as policies are enforced through the exchange of some trust-forming credentials. In a Grid environment, presumed trust is rarely feasible due to the dynamic and distributed nature of inter organizational relationships. Furthermore, trust establishment may be a one-time activity per session or it may be evaluated dynamically on every request. The dynamic nature of the Grid can in some cases make it impossible to establish trust relationships among sites prior to application execution.

## 2.2 Authorization Requirements

Authorization has several meanings in a Grid environment. For example: the process of issuing a proof of right, the proof of right itself (or reference to it), the process of determining an authorization decision by associating user attributes against access control policies, etc. Our interest focuses mainly on the third case because authorization is a key problem associated with the efficient function of the Grid environment, as authentication solely cannot effectively determine user rights.

Let us consider a situation where a multi organizational Grid shares resources. With current approaches, every change in any of the organizations' personnel requires the determination of new or changed user rights. This interaction places barriers and administrative overheads to the Grid functionality, as each participating organization acts as a resource provider or resource consumer. In such settings, expressing access control in terms of direct trust relationships between resource providers and consumers has the problems of scalability, flexibility, expressability and lack of policy hierarchy.

We can further classify the authorization information into two categories: (a) the general information regarding the user inside his own organization, like groups or roles he belongs to, along with other access control rights, and (b) information - permissions regarding what the user is allowed to do at the offered resources in the Grid. The first type of information should definitely be kept at the local users' organization, while the second one should be widely known or be able to determined by the rest of the participant organizations.

Finally, an authorization decision could be made either at the entry point of the shared resource, determining the user's access control rights, or, at a central point external to the shared' s resource organization. We argue that even though a central decision mechanism offers certain advantages, this should not overcome every organization's ability to decide for its own resources access permissions'. We also argue that these permissions could be very well modified during the organizations participation in the Grid and, furthermore, that an organization could decide to apply some granularity to these permissions and not to leave it open (same permissions) to anyone who wishes to use it.

## 2.3 Enforcement Mechanisms

The authorization problem can be further divided in two sub-problems. First to determine the permission set of a user and then to enforce access control rights by using existing access control mechanisms. The enforcement of access rights is usually done by determining and assigning the appropriate rights to the entitled user for the corresponding resources. In a typical non-distributed system scenario, the enforcement mechanism usually focuses on an application; because it is through this that a user actually accesses the underlying resource. The main concern here is the access rights of the user to that application and not the application's to the resource, as the application is a trusted software component and its access rights are easily determined.

In a Grid context however we anticipate another scenario. Resources are accessed by software components not necessarily trusted by the resource owners. The resources act as hosting environments for these software components (services), which are often transient and migrate between different resources to meet performance criteria. This implies that it is not possible to establish static trust relationships between the service and the underlying resource's hosting platform, while it is essential to examine the access rights of the user to the service but most important the access rights of the service to the current resource environment.

We can identify two categories of enforcement mechanisms: the application dependent and the application independent. Usually the application dependent mechanisms are directly integrated in the service and exercise access permissions before the service attempt to access the resource. The problem here is that the resource's hosting platform should trust the service's code and therefore it is favoured in situations where services are stationary. In the Grid environments however we believe that the application independent mechanisms are more favourable. In the later case the mechanisms are separate from the service's implementation.

# 3 CURRENT APPROACHES

Current Grid security technology is sufficient to address computational problems in healthcare. However Healthgrid technology is not restricted to the use of Grid technology for distributed computing only. Eventually, Healthgrids should offer a generic platform for all eHealth actors. Sharing of large amounts of distributed heterogeneous (on various

levels) data is therefore an important point of attention.

There are several examples of existing systems that attempt to address the authorization problem in a grid-computing environment. In the community authorization service (CAS) case for example, a CAS server is in charge of authorization of a community users, while resource providers in the community only need authentication after delegating authorization functions to the CAS server. In the generic authorization and access control (GAA) system, GAA API functions obtain policies from local files, distributed authorization servers and from credentials provided by the user. Its goal is to design a flexible and expressive mechanism for representing and evaluating these authorization policies. The Global Grid Forum is another system working on a grid security architecture standard (Lorch et all, 2003) which is based on Web Services security studies and tries to define fine-grained and coarse-grained authorization mechanisms under the Open Grid Service Architecture (OGSA) framework (Siebenlist et all, 2001). GEMSS is another approach that is based on a public key infrastructure, and implements end-to-end security mechanisms in line with the web services security specifications (Herveg et all, 2004).

Several security authorization models have also been proposed applying the RBAC access control model. Al-Kahtani and R.Sandhu, proposed for example a variation of the RBAC model, the rule-based RBAC (Al-Kahtani and R.Sandhu, 2002). According to their proposal Rule-based RBAC, rules are used to produce roles. These rules are expressed by attribute values, which the potential user presents to the system. These values have some seniority between them, which makes them comparable. The model has been used in cases of hosts visited by a huge number of visitors not known in advance by the system. Al-Kahtani also proposed a series of modified RBAC models which can very well contribute to the theoretical basis of our work. The PERMIS access control policy also provides a tool that takes advantage of the RBAC model by enabling roles to users. User's attribute certificates in the form of x.509 are compared against a list of permitted roles in an LDAP directory and a decision to grant or deny access is granted. An exact match is however still required in order to acquire a role (Chadwick, Otenko, 2002).

All these approaches are based however on certain attributes and predefined permissions that should be exactly fulfilled and thus they do not adequately handle the dynamic nature of the HealthGrid environments

# 4 A DYNAMIC AUTHORIZATION FRAMEWORK FOR HEALTHGRID ENVIRONMENTS

Membership in such a complex and distributed environment like the HealthGrid is dynamic as participant organizations may join or leave any time. Hence the relationships among participants are not constant while access control policy at every participant changes frequently. Authentication and role assignment therefore represent a major problem in HealthGrid environments. The following concepts are essential for the proposed authorization framework to be described later.

## 4.1 The Use of Mobile Agent Technology

The use of mobile agent technology can be useful in such environments. Mobile agents provide an undoubting advantage over the traditional client-server paradigm. The ability to move computations across the nodes of a wide area network helps to achieve the deployment of services and applications in a more flexible, dynamic and customisable way. Mobile Agents are often independent of particular hardware or operating system, and can be deployed in heterogeneous environments. Users belonging to different participants of the HealthGrid system try to access resources via the use of mobile agents. Once let loose, mobile agents roam the Grid network, seek information, carry out tasks on behalf of their senders autonomously, and return to present the results of their queries. A doctor for example working in a hospital A may seek information regarding a patient. This specific patient might have been treated in other medical institutions where the doctor obviously does not hold an account. Instead the doctor could send an agent loaded with credential which would help him to access a local security role at every hospital that visits and thus help him gather information regarding the specific patient.

However, agent technology carries with it associated security vulnerabilities that had to be addressed in order that this new technology to be of any use. A key requirement in the process is to find a flexible, convenient and effective approach to deal with the mobile agent authorization problem. The approach that solves the problem, dynamically maps an unknown user to predefined organizational roles based on the unknown user's credentials and role-

assignment mechanism. The agents should have proper authorization mechanisms for providing entry to the agents. The mechanism should automatically assign roles to agents. Our approach is based on the principle that when an agent migrates to a specific platform, the host decides what privileges to grant to the requested agent. Instead of using the traditional access control method via an access matrix, the mobile agent acquires a role from the local role hierarchy. This is a direct use of RBAC approach, but with a flexible method of selecting the appropriate role every time an agent arrives at a hosting platform.

The different hosts constituting the HealthGrid system are visited by mobile agents and should be able to automatically assign them different roles. This should mainly depend on a set of credentials (values on attributes) that the agent carries with it, the local security policy and constraints defined by the host that provides the resources. In any case the organization that provides the resource must have the privilege, the authority and flexibility to define the set of rules that define the security roles, applied to the specific organization.

## 4.2 The Role Decomposition Process

The process that decomposes the security roles is based on the following: the roles are defined using rules, which are constituted by Attribute_Expressions (A_E). These A_E are actually constructed by using Attribute_Fields (A_F) related to each other with some level of seniority as we shall see later on. In order to define a specific role, these A_Fs are filled with Attribute_Values (A_V), which are the ones that finally describe the role (fig 1).

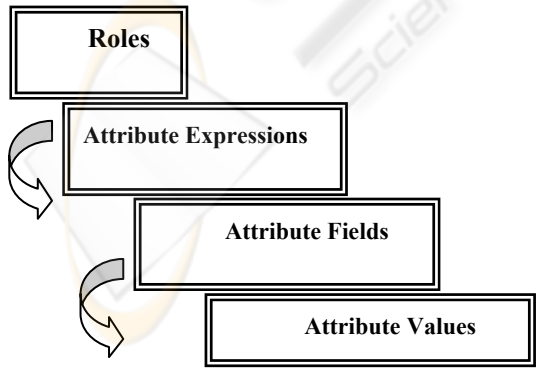


Figure 1.

The key component in the above process is the use of the A_Fs, because in order to be able to exchange agents, the different hosts should refer to common sets of credentials. This implies the existence of common A_Fs between organizational hosts, in order to define a common ground of understanding. This assumption is already a necessity in a Grid System environment as the different organization‘s scope is to facilitate resource sharing and problem solving.

The A_Vs that an agent carries with it might correspond to the A_Fs required for a specific role. In most cases however, they will not literally equal them. And in the case that these A_Vs are not of numeric type, a comparison between these two Attribute Expressions is not possible, even though they correspond to identical A_Fs. F. ex lets consider an A_F: “Business_Position” and a role requiring the A_V: Stuff Manager for this specific A_F. What happens if a mobile agent presents an A_V: General Manager? How could it be evaluated?

## 4.3 The Seniority Concept

A solution to the problem would be the introduction of the concept of seniority between the A_Vs of an A_Fs and furthermore seniority levels between A_Fs. The seniority level between A_Vs indicates which value dominates which in an actual comparison while seniority levels between A_Fs denotes which A_F consider higher in seniority to another A_F when it comes to compare A_Vs belonging to these two Fields. This leads us to the definition of the concept of an A_F hierarchy. The above mechanism permits a comparison between A_Es even though they are not constituted by exactly the same A_Fs. A_Es could be comparable as long as they either have identical structure (A_Fs), or A_Fs related each other with some level of seniority.

Instead of storing and distributing lists of A_Vs, which are difficult to manipulate, we introduce the idea of transforming the nonnumeric values to numeric. This makes the comparison between A_Vs extremely ease and self-explanatory. Moreover the A_Fs can contain not only the corresponding A_V but also the scale that this value came from. This is absolutely necessary, as different organizations might not use exactly the same scale to describe the A_Vs of the same A_F. Sending the scale with the value would permit them to “transform’ it to their own scale, performing sort of normalization. This is another degree of liberal mechanism in our model, as the main motivation remains always to permit every participant organization to establish its own security policy as independent as possible.

Obviously, this transformation is not working for all kind of A_Fs. There are certain categories of A_s which is meaningful to try to transform them into

numeric (e.g. names). In these cases where the A_Vs are discrete and determine, they should stay unchanged and left to the hosting organization to decide about the necessary seniority between these A_Vs (e.g countries, locations).

By this, A_Fs can contain only numeric A_Vs. The introduction of seniority between A_Fs and between A_Vs eases the comparison between A_Es because:

1. It imposes seniority between A_Fs, while it not even necessarily agreed in advance between all participants of the infrastructure. Instead every participant reasons and evaluates the importance of each A_F under his/her own judgment.
2. A_Fs are grouped and related with some level of seniority according to their significance. Not all A_Fs necessarily are related each other. Introducing seniority between A_Fs with diverse implication is meaningless, so only A_Fs with similar significance are related to each other. This leads us to groups of A_Fs with seniority among them.

A role can be expressed in the form of an A_E as follows:

| | | | | |
|---|---|---|---|---|
| $X_{11}$ | $X_{12}$ | $X_{13}$ | $X_{21}$ | $X_{21}$ |

where:

$X_{ij}$ represents an A_F, i identifies the group it belongs, j represents the rank number of the A_F inside the group while $X_{ij'}$ represents the corresponding A_V.

The above mechanism allows for two roles from different role hierarchies R and R' to be compared for dominance.

# 5 THE ADAPTED MOBILE AGENT STRUCTURE

The effective management of trust and policy within a community in HealthGrid systems also requires flexible, autonomous mechanisms able to take into account, when organizing communities, not only the semantics of policy statements, but also the ability to negotiate policy terms and to manage restricted delegation of rights. A procedure that is already known in agent technology.

The mechanism presented earlier could therefore be implemented easily using existing agent technology. In this section we present such an adaptation of agent technologies in order to enforce a flexible and scalable authorization service, suitable for complex HealthGrid system structures. We assume that the authorization decisions must often be made in the absence of strong existing trust relationships. Furthermore we assume that the proposed authorization service will use the existing access control model at every participant node of the Healthgrid. A suitable, flexible agent structure is needed to anticipate the above requirements. The proposed adapted mobile agent structure is based on the following concepts:

When an author (programmer) constructs an agent, for such an environment, he should therefore include in this case the following:

- The source code
- The set of possible A_Fs, the agent is allowed to carry
- Default A_Vs, for the corresponding A_Fs which construct an A_E that yields to a default agent role.

Accordingly, when a user is about to dispatch an agent, he should include:

- All the previous, plus
- A set of A_Vs that correspond to the set or subset of the A_Fs the author of the agent, permits it to carry. These A_Vs derive from the role that the user is already been assigned at the specific host.

Every mobile agent has a default agent role defined by its author (programmer) and a user role inherited by its user/sender. The corresponding A_Vs do not change during the agent's journey, as they constitute its original identity, authority and credibility.

Let us consider the generic case where an agent initially launches on host H0 holding its default role R0. The user (sender) loads his role $R_U$ on the agent thus the A_Vs that actually constitute this role. We assume that the agent will execute autonomously on a set of network hosts (namely $H_1$, $H_2$ …$H_n$). When agent is residing on the host $H_i$, its current role is $R_i$. We have to distinguish the following different role symbolisms:

- $R_0$ stands for the agent's default role
- $R_U$ stands for the user (sender) role that is initially assigned to the agent
- $AR_i$ stands for the agent role, the role which the agent holds while migrating from host $H_i$ to $H_{i+1}$,
- $R_i$ stands for the role that the agent acquires on host $H_i$
- $R_{i+1}$ stands for the final role that host $H_{i+1}$ grants to the agent.

This ultimate role assignment depends on the role of the arriving agent $AR_i$, the specific migration method the agent followed to arrive at host $H_{i+1}$ and the agent's initial $R_U$ and default role $R_0$. This means that when an agent migrates from host $H_i$ to the next host $H_{i+1}$, its role (represented as $AR_i$) is actually the role $R_i$ that the agent obtained during its execution

on host $H_i$. According to (Lorch et all, 2003), the agent migration to the next host can be initiated by two different ways: either from the hosting place (host), or from the agent itself, (the agent's code decides to move to the next place). The migration can also take place with two different ways: either by handoff (a principal hand his authority to another principal), or by delegation (a principal is combining his authority with another principal).

During the process of delegation, when one principal (initiator) authorizes another principal (delegate), the attached privileges, in the form of A_Vs, are passed from initiator to the delegate. Te proposed delegation mechanism supports the RBAC access control by delegating roles and provides a higher level of granularity than approaches limited only to individuals. The proposed mechanism can be summarised as follows:

Every host in the Grid is considered as a unique entity and is able to delegate his authority. The four distinct cases (namely the specific migration method as well as the initiative) and our proposed role decomposition methodology lead us to:

*1. Place handoff:*

The current host $H_i$ hands off his authority to the next host $H_{i+1}$. In this case the next role $R_{i+1}$ will result from the following expression:

$R_{i+1} \leftarrow AR_i \Rightarrow R_{i+1} \leftarrow R_i$

The A_Vs the agent came with at the host $H_i$ are presented and at the next host $H_{i+1}$

*2. Place delegation:*

Host $H_i$ delegates his authority to $H_{i+1}$. In this case the next role $R_{i+1}$ will result from the following union expression:

$R_{i+1} \leftarrow AR_i \Rightarrow R_{i+1} \leftarrow R_i \cup AR_{i-1}$,

When the host delegates his authority to a visited mobile agent, the A_Vs that constitute the role that the host assigned to the agent, are transferred to the delegated agent. The A_Vs of the roles $R_i$ the agent acquired on host $H_i$ are combined with the ones that came with from the previous host $H_{i-1}$, in order to obtain the next role $R_{i+1}$.

*3. Agent handoff:*

The agent directly hands off to $H_{i+1}$. The next role $R_{i+1}$ will result from the user's initial role, the role assigned by the agent's user. The following expression stands in this case:

$R_{i+1} \leftarrow AR_i \Rightarrow R_{i+1} \leftarrow R_u$

The A_Vs of role $R_u$ are presented to host $H_{i+1}$ in order to obtain the next role $R_{i+1}$.

*4. Agent delegation:*

The agent can delegate itself to $H_{i+1}$. Since the agent delegates his authority, his default role $R_0$ is combined with the agent's role $R_U$. The next role $R_{i+1}$ will result from the following expression:

$R_{i+1} \leftarrow AR_i \Rightarrow R_{i+1} \leftarrow R_U \cup R_0$

The A_Vs of roles $R_U$ and $R_0$ are combined in order to obtain the next role $R_{i+1}$.

The above agent migration methods affect decisively the agent's ability to join the security policy of the different Healthgrid organisations it visits. This adapted agent structure can therefore effectively address the authorization problem in HealthGrid environments and thus facilitates collaboration between participant health institutions in order to establish an e-health environment. .

# 6 APPLICATION – RESULTS

The proposed process and the related assignment – migration algorithms presented above have been implemented and applied, using the Java programming language. The implementation used as an example a typical health care environment where the ministry, local hospitals, hospital departments (nutrition, etc), patients and health insurance companies collaborate in order to exchange information. We tried different role hierarchies at every organization and we tested all different migration methods. Every time the mobile agent adopted a different suitable security role at the organization that it visited.

The results have shown that the proposed framework is applicable and implementable, and can be applied in real life situations. Also the proposed framework does not require any change of the different local role hierarchies of the participating organizations and it preserves their security requirements. Based on those results, it can therefore be said that the experimental implementation of the proposed adapted agent structure, together with the proposed earlier role assignment process, can effectively address the authorization / security problem in Grid based health environments

# 7 CONCLUSION AND FUTURE WORK

Securing HealthGrid environments is a major issue today. Current Grid security technology is sufficient to address computational problems in healthcare. Healthgrids are not restricted to the use of Grid technology for distributed computing only. Healthgrid should eventually offer a generic platform for all eHealth actors. Sharing of large amounts of distributed heterogeneous (on various

levels) data is therefore an important point of attention. In this paper we discussed the overall security problem of such environments and we presented a suitable a local authorization policy for HealthGrid participating organizations, which utilizes the advantages offered by both the grid and the agent systems technology. More specifically, we addressed the dynamic authorization problem of HealthGrid environments by describing a flexible, RBAC based role assignment mechanism. Our approach proposes a secure and consistent solution to the authorization problem in HealthGrid environments, based on a role decomposition process, that every organization is qualified to perform according to its local security policy. The main contribution of this work is that it proposes a practical mechanism by applying the well-accepted RBAC access control model to dynamic HealthGrid based environments. The proposed methodology has been successfully tested in a real life health environment.

## REFERENCES


Al-Kahtani, R.Sandhu, 2002 A model for Attribute-Based User-Role Assignment. *Proceedings of the 18th Annual Computer Security Applications Conference*, Las Vegas.

Chadwick, Otenko, 2002. "The PERMIS X.509 Role Based Privilege Management Infrastructure", *Proceedings of the seventh ACM symposium on Access control models and technologies SACMAT 2002*, Monterey, California, USA

Ferraiolo, Sandhu, Gavrila, Kuhn, Chandramouli: 2001 Proposed NIST standard for role-based access control. *TISSEC 4(3): 224-274.*

Foster I, Jennings N, Kesselman C, Brain Meets Brawn: 2004 Why grid and Agents Need Each other, *AAMAS'04, NY, ACM*.

Herveg, J., Crazzolara, F., Middleton, S. E., Marvin, D. J. and Poullet, Y. 2004, "GEMSS: Privacy and security for a Medical Grid". *In Proceedings of HealthGRID 2004,* Clermont-Ferrand, France.

Jennings, N., 2001.An agent-based approach for building complex software systems. *Communications of the ACM, 44(4). 35-41*.

Lorch, B.Cowles, R.Baker, L,Gommans, P.Madsen, A.McNab, L. Ramakrishnan, K.Sankar, D.Skow, M. Thomson, 2004 Conceptual Grid Authorization Framework and Classification, *Global Grid Forum*, ..

OASIS 2003, Security Services Technical Committee XAMCL, extendible access control markup language (XACML) *committee specification 1.0.*

SAML 2003 Security Services Technical Committee, Assertions and protocol for the oasis security assertion markup language (SAML), *OASIS* .

Scott Richard., Jennett Penny, Yeo Maryann, 2004, Access and authorisation in a Glocal e-Health Policy context, *International Journal of Medical Informatics* Elsevier (2004) 73, 259—266

Siebenlist, V. Welch, S. Tuecke, I. Foster, N. Nagaratnam, P. Janson, J. Dayka, A. Nadalin, 2002, OGSA security roadmap, Open Grid Services Architecture Group.

Simpson Andrew, Power David, Slaymaker Mark, 2006, On tracker attacks in health grids, *SAC'06* April 2006, Dijon, France

Welch, F. Siebenlist, I. Foster, J. Bresnahan, K. Czajkowski, J. Gawor, C. Kesselman, S. Meder, L. Pearlman, S. Tuecke, 2003. Security for Grid Services. *Twelfth International Symposium on High Performance Distributed Computing (HPDC-12*), IEEE Press