

A policy-ruled Knowledge Dissemination Architecture for supporting Multi-Domain Secure Interoperation

Stefanos Gritzalis

Laboratory of Information & Communication Systems Security (*Info-Sec-Lab*),
Department of Information & Communication Systems Engineering,
University of the Aegean, Samos, GR-83200, Greece
E-mail: sgritz@aegean.gr URL: www.icsd.aegean.gr/sgritz

Abstract

The systematic and effective management of existing knowledge can significantly assist organizations develop quick mechanisms able to respond to various change-driven demands. Knowledge Management (KM) systems have been effectively deployed for more than two decades, incorporating Information Retrieval techniques as well as tacit knowledge utilization. However, traditional KM systems have limited functionality, mainly because of their adaptation to the single-organizational model. When applied to distributed, intra-organizational coalition framework, consisting of cooperating organizations, they can significantly leverage the overall organization's performance. In this paper we present the extensions to the existing security models that must be realized prior to the employment of a distributed knowledge-exploiting architecture; moreover, heterogeneity, interoperability and scalability related issues are considered in order to handle the diverse set of platforms found in different Information Systems. We also propose a distributed KM framework for secure interoperation, focusing on the technical challenges and the available solutions.

Keywords

Security, Distributed Knowledge Management

1. Introduction

The advent of networked infrastructures, as well as the exploding expansion of digital repositories, urge towards Information Systems redesign such that they support efficient and accurate knowledge retrieval. The diversity of existing platforms even within a single organization, has been the main obstacle as far as knowledge exploitation is concerned. Knowledge Management (KM) systems attempt to mitigate these problems by providing a cohesive set of solutions that attempt to overcome heterogeneity problems, as well as to facilitate knowledge sharing or tacit to explicit knowledge transformation.

For more than two decades KM systems are utilized by several organizations, sometimes supporting their core business functions and thus consisting among their competitive advantages. In addition to Information Retrieval, which is the core functionality of a KM system, the real impetus to an organization is given through the effective support of the socialization process that enables the utilization of the intellectual capital, namely the human factor. Even in the early frameworks of organizational theory the support of inter-organizational knowledge exchange emerges as a real challenge.

Traditional KM systems fail to utilize knowledge that resides in other organizations. Their centralized orientation and their closed architecture, prevent from exploiting organization's intellectual capital to its full extent. The advent of networked infrastructures on the other hand, arises as a challenge for cooperation between Information Systems, in terms of knowledge exchange and diffusion. Even though distributed architectures in KM emerge as a promising solution, they pose a number of serious problems [BON02], concerning the

efficient and secure management of jointly owned resources. The complexity of the problem mainly resides in the contradictory requirements of such a task: on one hand there is a necessity for sharing and collaborative contributing to the knowledge assets, while on the other hand it is necessary to protect knowledge assets so that they will be accessed only by those who retain legal access rights.

In this paper we address the security issues related to the management of knowledge residing in distributed collaborating autonomous domains and specifically how to ensure that the knowledge assets can only be accessed by “roles” - from different organizations - that retain the legal rights to do so. We present the key challenges in the sector of multi-domain environments knowledge exploitation, on the grounds of utilizing three core technologies: software agents, ontologies and security policy languages. Our approach allows determination of access control on the ground of XML-based security policy languages, while it is characterized by its extended scalability potential. Furthermore, through the use of software agents it enables transparent identification of knowledge assets, as well as semantically enabled query formation. The rest of the paper is organized as follows: in Section 2 we argue about the access control related challenges and the advantages of using security policy languages, Section 3 presents the main principles of the Role Based Access Control (RBAC) model and the necessary adaptations for multi-organizational environment, Section 4 highlights the related work while Section 5 presents the technological solutions that are adopted by our approach.

2. Access Control Solutions for Distributed Environments

The rapid growth of Information Systems (IS) and the emergence of high performance networked systems did not come without its drawbacks. Security breaches are a daily phenomenon, caused by outside intruders as well as by insiders. Managing the resources of a distributed system is a big challenge that requires a lot of effort on both the design as well as on the implementation of countermeasures. Security policies are adopted to a high extent towards this direction [LUP99].

A typical enterprise system includes a large number of heterogeneous devices, which run a variety of applications and offer services to a large number of users with diverse authoritative rights over these resources. The complexity of the managed systems reflects to a demand for high costs in both terms of time, humans involved and long development cycles. Moreover, these requirements become critical for a number of reasons: (i) management must be *distributed* in order to be scalable and cope with the size of enterprise networks, and (ii) management procedures must be *automated* to reduce administrative costs. Networked environments are designed to be highly adaptable to support rapid deployment of customized services. Thus, management also needs to be dynamic and flexible to deal with the evolution of the systems being managed [DAM00].

The aforementioned identified requirements for management systems can be facilitated with policy-based management approaches where the support for distribution, automation and dynamic adaptation of the behavior of the managed system is achieved by using security policies. The main benefits from using policies are improved scalability and flexibility for the management system. *Scalability* is improved by uniformly applying the same policy to large sets of devices and objects, while *flexibility* is achieved by separating the policy from the implementation of the managed system. Policy can be changed dynamically, thus changing the behavior and strategy of a system, without modifying its implementation or interrupting its operation. Policy-based management is largely supported by standards organizations such as the Internet Engineering Task Force (IETF) and the Distributed Management Task Force (DMTF), and most network equipment vendors.

As enterprises are increasingly leveraging Internet technologies to adopt e-business practices, they expose internal resources to customers and require that enterprise-wide authorization policies are easily established and implemented. Authorizations must be

enforced both at the application level and in network elements, and must be explicit, i.e. authorization policies must be precisely and unambiguously stated to define the set of acceptable requests. Various techniques have emerged for programming network elements to support adaptive services, such as active networks, mobile agents, and management by delegation. While these approaches support the programming of new functionality into network elements and host devices, they increase the security concerns regarding access to network resources and services. Authorization policies must therefore specify which users are permitted to program network elements, which services users are permitted to access and under what circumstances.

The proliferation of non-integrated security mechanisms and products, each with independent administration and application development interfaces, leads to separate authorization policy implementations within each individual application and system. This makes it difficult to support global security policies in accordance with enterprise access control goals. It should be possible to provide consistent security across the distributed object system and associated legacy systems.

3. RBAC and Adaptation to Collaborating Environments

Policy must be separated from the security mechanisms that enforce access control in order to enable the specification and integrated administration of global policies. Several challenges arise on this field, due to the very large number of subjects - resources that need to be administered and also due to the very large number of users. The Role Based Access Control (RBAC) [SAN00] model seems to be dominant and widely accepted both in academic and commercial environments. The main principle of RBAC is related with the fact that normally users with similar roles need to be accredited for the same actions and need to have the same access rights. By classifying users to roles and accordingly by associating individuals with a role, the security management is simplified dramatically. For example, each time somebody enters the organization, we simply classify him/her to one of the predefined roles. Accordingly, when somebody leaves the organization, we do not need to manually withdraw all the access rights for every resource she was assigned to have access rights.

The basic concepts behind RBAC definitions are *users*, *roles*, *permissions* and *sessions* (Fig. 1). Users are the physical entities and they can be classified to roles according to the tasks assigned within the organization they work for, while privileges are assigned to roles, in order to enable them perform certain tasks by accessing the necessary system's resources. Sessions serve the case where a user signs on a system to perform some specific task. They introduce some abstraction between users and roles and thus enable users to activate temporarily different roles, or to log with two or more roles at the same time.

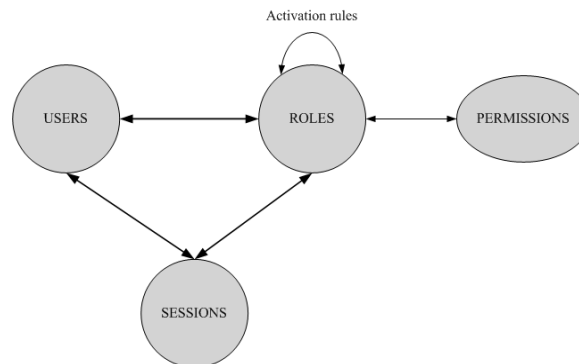


Fig. 1 Basic RBAC concepts

RBAC supports “many-permissions to many-role” and “many-users to many-roles” assignments. Embedded in the standardized RBAC is the support for cardinalities that limit the number of users able to activate a role. The standardized RBAC covers more of the

demands for access control enforcement, still though for dynamic distributed environments (or as in our case multi-domain environments) demand several modifications. Extensions to the basic RBAC model include incorporation of role-specific parameters and association between them. For example time periodicity that determines the time activation limits for a specific role to be active. This adds extra complexity when defining associations between role relationships. Other supplementary extensions are associated with context related predicates in role definition especially in multi-domain environments (for example the domain where a specific user belongs to, personal data related parameters, etc). In general the RBAC model has proved to be dominant due to its flexibility and scalability support. Therefore, in order to apply role based authorization schemes to multi-domain environments a considerable number of extensions are required. Such modifications have been accounted for by our approach.

In our role specification we consider attributes that:

- Are pre-settled by the system administrator and that the user must specify in order to be granted authorization to activate a role (mainly these are domain specific parameters, correlating a user with a specific domain. For example all users who belong to domain `oncology.genhospital.org` or to the domain `hematology.genhospital.org`)
- Enable time periodicity, for example allow access within pre-specified time-intervals
- Enable context based role-assignment and role correlation for roles specified in different domains.

3.1 Multi-domain Environments

Security policies have satisfactorily resolved many problems in distributed environments; though, they have been applied mainly to the single-organization paradigm. A more complicated situation is related with the attempt to create a policy-managed collaboration scheme between different organizations. In most cases, establishing a collaboration access scheme involves off-line negotiation, as well as complex procedures like identification of the negotiating parties, agreement upon the conditions for sharing a resource etc.

Two system types can be considered under this collaborating framework: *peer-to-peer networks*, and *autonomous domains*. Peer to peer networks resemble communities with common interests, the terms of bounding though are more loosely coupled than autonomous domains. The second category of systems can be met in many real-life systems, such as e-government environments, or healthcare systems, which consist of several cooperating hospitals. In the latter case the sensitivity of the data poses more security restrictions and thus establishing a common state for knowledge exchange requires that organizational roles are well defined in terms of access rights and obligations based on the grounds of a well-stated security policy, while a common access state between different organizations is unambiguously allocated.

We can classify these systems, according to the access models they adopt, to the following two categories:

- *Trust based systems*. The notion of trust is introduced mainly in complex, non-hierarchical or inter-related systems such as the Internet, where totally unknown roles might be interested to enter into relations between them, or to cooperate on financial terms basis. This situation is common on Internet transaction systems, such as e-commerce etc. The authorization of a transaction is based on the basis of estimating the cost and the substantial loss for a specific role
- *Autonomous systems*, with a well-formed security policy and a well-defined organizational structure.

We will restrict our scenario to the second system category, which is characterized by well-defined organizational policies and by the fact that the systems cooperate on the basis of a commonly agreed target, such as improving the efficiency of the governing infrastructure or minimizing the response time for the treatment of patients within the national healthcare system.

3.2 Role Mapping Across Domains

One of the main problems in our approach is related with the concurrent treatment of security parameters between the cooperating systems. For example, one person could be appointed to work as a manager on one system and also as a restricted user on another. In order to be able to perform the necessary actions on each system, there has to be a way to enable him/her to provide his/her credentials and to be granted access on both systems, no matter of the location. A basic obstacle is related with the fact that different organizations have different organizational schemes -roles in general - and also different privileges assigned to each role. We will confine our selves to the RBAC model and we will assume that all the participating domains adopt this model for access control. Trying to merge different access control schemes can be risky, since it's not exactly clear how RBAC roles could be correlated with a system following different principles for privilege assignment for example Discretionary Access Control (DAC). Another issue also arising is how it would be possible to automatically enable user-authentication between different domains by merging the system's different policies. The subject of automating a negotiation procedure that would result in performing automatically a matching between the roles, their security attributes and their mappings can be proved to be NP-complete [BHA03]. Many approaches have been proposed, others semi-automated [GLI01], others focusing on providing a scalable and efficient solution, even with human interventions to some extent. We believe that in most of the real scenarios the sensitivity related with protection of data, as well as the legislative restrictions make the automated negotiation procedure - due to the consequences in case of wrong mismatching between roles - an endangered solution and therefore other approaches, less error prone should be adopted in most cases. Therefore, we introduce the following solution: we create a role hierarchy, to which the local role representation schemes must conform, and the administrators who are aware of the legal and technical implications of the mapping are responsible to perform the correct mapping between the collaborating domains. In our case, coalition administrators will produce a role mapping between the different organizational policies, through the general intermediating mapping scheme. This mapping scheme acts as a federal, and all the domains wishing to cooperate need to comply with it (Fig.2). This solution can be considered as having the following characteristics:

- Robustness, due to the fact that each user can be assigned only the roles that have been defined for him/her by the coalition administrators. Users cannot gain unauthorized access since under the multi-domain environment, only pre-specified mappings are allowed. In order to enforce authorization decisions, the users attributes can be evaluated and in case of sensitive environments, the use of asymmetric key cryptography techniques may be necessary as an extra security measure. In this case, the user might be asked to provide his/her key with every access request in order to verify the authenticity and origin of the request. Encryption of messages may also be employed, whenever it is necessary to transmit sensitive messages over non-trusted networks paths.
- Scalability, because the way mappings are defined enables the system to grow without adding extra costs and without raising the system's complexity management.




Fig. 2 Role-mapping across different domains. The role X on the federal role hierarchy scheme maps on role Y on domain's B hierarchy and role Z on domain's C role hierarchy scheme.

4. Related Work

Several frameworks have been proposed in recent literature, mainly utilizing peer-to-peer technology for sharing knowledge resources.

XAROP [TEM04] is a peer-to-peer system, which manages heterogeneous knowledge sources by using ontologies. The security model adopted by this approach relies on the use of public key infrastructure techniques, managed by a supervising entity specified within the XAROP infrastructure. Though, the determination of access privileges is not defined in a flexible manner. Classification of documents requires by the user to manually classify users and groups, putting the main security management burden to each user. Furthermore, it raises serious problems when it comes to manage scalability related requirements. Finally, storing access permission details for each document within a distributed system becomes an untreatable problem.

ADAM [SEL04] is a distributed system, which utilizes trust based negotiation procedures for the establishment of transactions between users. Its architecture is agent based, with one agent being responsible for gathering the knowledge from distributed nodes and a second agent for handling the authorization processes on behalf of the user. ADAM mainly manages knowledge related to its users and not its consisting assets; the authorization process is mainly based on the notion of trust. Trust based systems collect information about the participating parties in a certain transaction prior to allowing authorizations, by questioning other users or entities and accordingly they evaluate the request based on the user's previous reputation. Trust models are suitable for environments not hierarchical without a well-defined organizational policy, like for example the Internet. Internet transactions, e-commerce environments are suitable for trust based security models. In e-commerce for example, a misclassification of a malicious user as reliable may result in potential financial loss. In case of critical environments though, the deployment of reliable mechanisms (which reason for a specific request by examining the predefined policy and by examining the user's credentials) is necessary in order to avoid security breaches.

SemanticLIFE [WEI04] is a project that stores an individual's entire digital life and makes it available to coworkers. The security scheme effectuation is mainly based on implementing role-based access controls through the usage of database systems. There is no support for intra-organizational collaboration, nor this access control model provides support for cooperation between different systems, while its scalability potential can be also put under consideration. Security policies, as utilized in our architecture, offer a more flexible and much more robust way for security administration, while they can provide the system with greater scalability potential without lacking in attack resistant metrics.

The aforementioned systems mainly aim at utilizing knowledge and making it available to all community members. They do not provide role-specific authorization, nor they support a robust and scalable solution for multi-domain authorization.

Our system provides solutions for the following tasks:

- *Knowledge assets discovery*, which is handled by security agents and is based on the use of the appropriate ontology
- *Authorization process*, which investigates if the user that requests access to an asset has the appropriate level of classification
- *Negotiation*, when it comes to inter-organizational knowledge transfer.

By deploying security policies, we enable the participating domains to maintain their autonomous character, while with the proposed solutions for policy mapping and cross-organizational role assignment we enable a robust and scalable solution that enables knowledge dissemination in a flexible and secure way.

5. Technological Framework Adopted

5.1 Policy Language

In our approach we utilize the Extensible Access Control Markup Language (XACML) [XACML]. XACML is a policy language that supports prohibitions, obligations, and resolution of conflicts. Its expressiveness and XML (Extensible Markup Language) codification support allow its integration on a variety of environments, such as web-service based environments, distributed autonomous systems, and with some modifications to be applied also to pervasive environments. Among XACML's strong points, are:

- It is standardized and it is open, allowing extensions that enable interoperation between various platforms
- It is codified in XML, which tends to dominate as codification standard while it is operating system independent.
- It allows extensions as to support the needs for a variety of environments.
- It allows context based authorization, which is a big advantage

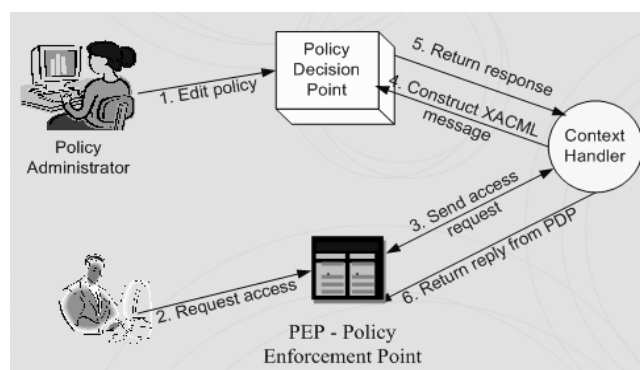


Fig. 3 XACML's overview

The XACML operates as follows (Fig.3): First the administrator is editing the policy and makes it available to the Policy Decision Point (PDP). Every request is directed to the Policy Enforcement Point (PEP). The PEP forwards the request to the PDP, which collects context-related information through the context handler. In addition we introduce time-periodicity related attributes to the parameters the context handler collects, while in the cases the

requester is making the request from another domain - through the policy mapping - his/her role is adjusted to the permissions of the corresponding role on the domain where the information is stored. The PDP then acquires all context information, the corresponding role for the new domain and loads the local policy and evaluates the request against the policy. The default policy considered in our system is prohibition.

The aforementioned policy enforcement framework is mostly useful for single domain authentication and authorization. For multi-domain environments, there is a necessity for multiple PDP and PEP deployment, one for each domain. In these cases, the authorization process works as follows (Fig. 4). Each domain maintains its own PEP and PDP, which reason about requests attempting to access resources within their range of responsibility. In addition, each domain maintains a set of mappings stored in a specific purpose registry maintained by the coalition's administrators and stored at the same server with the PDP module. Each request is directed to the domain's PEP. In case both the request and the requester belong to the same domain, the process works as described in Fig. 3. In case of a request for resources from remote domains, then the local PEP is constructing a message to the remote PEP. The request is directed to the remote PDP where it is evaluated according to the remote policy and the pre-defined mappings. In case an appropriate mapping does not exist, the request is not satisfied. The context handler facilitates the authorization process by enabling authorization based on attributes characterizing a user (or sets of users) from the requesting domain. For example, the remote context handler can enable remote authorization for all the members of the requesting domain, by examining the content of the request message. In a more detailed example, if the requesting domain is *oncology.genhospital.org*, the context handler by examining this domain specific information common for all users of this domain, it can authorize all of them to access specific resources; therefore, the authorization process can be simplified and performed faster.

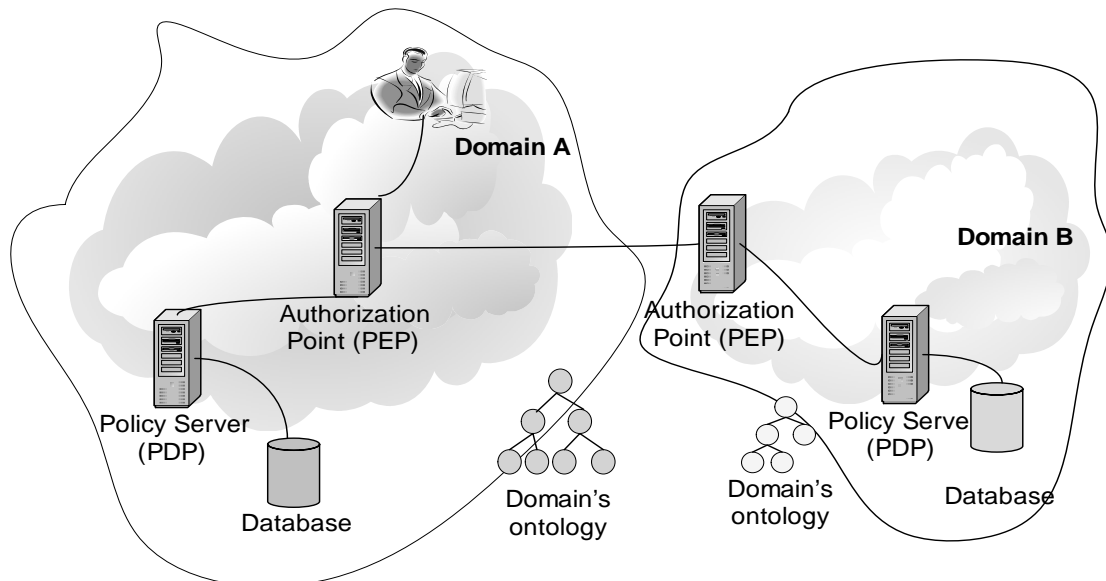


Fig. 4 Multi-domain authorization and access control enforcement

5.2 Agent and Ontology Technology

In order to facilitate searching and also to overcome heterogeneity related obstacles, the role of ontology can become crucial. Each domain maintains its own domain-specific ontology. We also introduce a central ontology repository, accessible from all the domains for retrieving domain ontologies. Ontologies define the concepts for each domain and their properties and enable semantic description of knowledge assets. Another requirement in our

research effort is to enable transparent knowledge assets identification. For this purpose we utilize software agents, which act on the user's behalf and query the distributed domains. Ontologies except from the aforementioned facilitation in semantic discrimination of knowledge assets, they also play a crucial role in facilitating communication between software agents as they enable standardization of terminology in agent communication messages.

In order to assist ontology-based knowledge discovery facilitation we introduce two concepts:

- The *Local Knowledge Repository (LKR)* represents the knowledge assets that are available in an organization's administrative domain. The repository might be a centralized one or distributed among servers and workstations of the domain. Both internal and external users of the organization gain access to its contents by enforcing local authorization policies. Knowledge assets can be stored in relational databases, XML files or any other type of data repository. In order to address heterogeneity issues that arise, since knowledge assets can be text, images, e-mails or any other kind of resource, LKR also contains a set of metadata, expressed in XML format, for asset description. Part of this metadata is a unique asset identifier, a short description in text, the author's name and a set of keywords used for matching user requests to knowledge assets.
- The *Global Ontology Repository (GOR)* that is accessible from any organization and is responsible for the storage and retrieval of domain ontologies.

The process of semantically enriched asset identification works as follows: For each domain a set of software agents is introduced: The knowledge discovery agent and the authorization agent. The discovery agent is responsible for identifying assets matching the user's query, while the authorization agent carries the user's credentials and facilitates the authorization process. The deployment of software agents enables transparency to the user, which poses a query to identify relevant to this query assets belonging to his/her domain or to a remote domain. Upon a specific query, the discovery agent accesses the Global repository retrieves and queries the domain ontologies. Each domain maintains a specific ontology, containing a classification of domain's assets according to their subject. The benefits from this process are two-fold: first we minimize the network resource consumption since we avoid posing irrelevant queries to the domains that do not contain topics of interest to the user. (This is also important in case where the network infrastructure is susceptible to instability due to hardware limitations, such as wireless networks) [MAL05]. Second, we also enable semantically enhanced asset identification for each domain's assets, as well as we enable treatment of heterogeneous files such as text, image files and so on.

5.3. Prototype implementation details

We present the basic features of a prototype that incorporates the aforementioned technologies and extensions to security models (Fig.5). Our prototype implementation consists of an organizational memory, consisting of the organization's past experience codified in semi-structured documents [BEL04a], while support is provided for several heterogeneous types of files, such as images [BEL04b]. The repository is implemented in Oracle 9i while Java technology is utilized for access purposes. Local and Global ontology repositories are maintained on each node; the local ontologies are facilitating semantic search for the domain, while the global ontology repository contains a summarization of the main thematic categories and facilitates the retrieval of the most relevant local ontology; therefore it directs the search to the most relevant domains. This architecture is deployed in different domains, each one maintaining its own autonomy. For each domain there is a policy decision point (PDP) which directs the policy enforcement point (PEP) to provide - or not - access to distributed resources of the system upon a user's request [BEL05].

Imagine the following system usage scenario. A user poses a specific query to identify topics and knowledge assets of specific interest. The discovery agent queries initially the local document management module. Accordingly it accesses the GOR and retrieves domain ontologies from which it identifies the relevancy of the user’s query with the repositories of specific domains. The messages exchanged between the domain specific agents are based on FIPA protocols, and the content embodied is based on the RDF ontology, which plays also a key role relative to the facilitation of heterogeneous assets knowledge discovery. For assets from the same domain, the authorization process as depicted in Fig. 3 is handled. The authorization agent carries the user credentials and directs the request to the PEP and the process will proceed as described in previous section. For all the assets belonging to remote domains as identified by the discovery agent, the authorization agent is forwarding the user credentials and the request to the remote PEP which forwards the request to its domain’s specific PDP and evaluates the remote request by evaluating its domain specific policy and the policy mappings as they have been defined by the coalition administrators. In case of adequate privileges, the user is granted permission and the discovery agent brings the requested assets to the user. Therefore, the user is provided with the chance to utilize knowledge from multiple domains transparently, where all the asset discovery processes as well as the authorization processes between different domains are treated by the system, through the use of the pair of Auth-Agent and Discovery-Agent, assigned to each domain. Another characteristic of our approach is that it enables context-based authorization based on the collection of user related attributes.

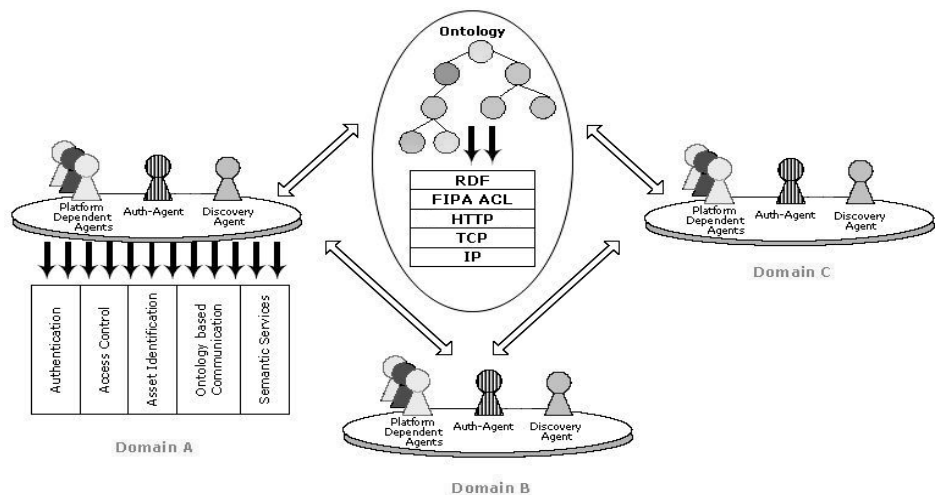


Fig. 5. Overall system architecture. The Auth-Agent and Discovery-Agent perform user authorization and asset identification respectively. We can also distinguish platform dependent agents.

6. Conclusions

Knowledge is a critical asset that can give competitive advantage to organizations, when accessed on time. The high volumes of data collected by organizations demand innovative and efficient methods for knowledge acquisition and dissemination. A big challenge is related with the ability to fetch and disseminate knowledge between different organizations considering the access control privileges associated with each role within each organization. Security policies can assist towards this direction, enabling determination of access control rights according to organizational roles; still, there is a demand for providing extensions to the traditional role based access control models. We have provided with a scalable and robust solution towards this direction, while we have also utilized agent and ontology technologies in order to achieve transparency and semantically enhanced knowledge retrieval facilitation.

We plan to continue our research work towards the establishment of semi-automated mappings between different roles and coalition formation.

7. References

- [BON02] Bonifacio M., Bouquet P., and Traverso P., “Enabling distributed knowledge management. Managerial and technological implications”, *Informatik – Informatique*, vol.1, 2002
- [LUP99] Lupu E., Sloman M. “Conflicts in Policy Based Distributed Systems Management”, *IEEE Transactions on Software Engineering*, Vol. 25, No 6, 1999
- [DAM00] Damianou N., Dulay N., Lupu E., and Sloman M., “Managing Security in Object-based Distributed Systems using Ponder”, in *Proceedings of the 6th Open European Summer School (Eunice 2000)*, Enchede, The Netherlands, 2000
- [SAN00] Sandhu R., Ferraiolo D., and Kuhn R., “The NIST model for role-based access control: towards a unified standard”, in *Proceedings of the 5th ACM Workshop on Role-Based Access Control (RBAC’00)*, pp. 47–63, 2000
- [BHA03] Bharadwaj V., Baras J., “Towards automated negotiation of access control policies”, in *Proceedings of 3rd IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY’03)*, 2003
- [GLI01] Gligor V. D., Khurana H., Koleva R. K., Bharadwaj V. G., and Baras J. S., “On the negotiation of access control policies”, in *Proceedings of the 9th International Security Protocols Workshop*, Cambridge U.K., LNCS 2467 Springer, pp. 188–201, 2001
- [TEM04] Tempich C., Ehrig M., Fluit C., Haase P., Marti E.L., Plechawski M., Staab S. “XAROP: A Midterm Report on Introducing a Decentralized Semantics based Application”, in *Proceedings of Practical Aspects of Knowledge Management (PAKM 2004)*, Vienna Austria, LNAI vol. 3336 Springer, pp. 259-270, 2004
- [SEL04] Seleznyov A., Mohamed A., Hailes S. “ADAM: An agent-based Middleware Architecture for Distributed Access Control” in *Proceedings of the 22nd International Multi-Conference on Applied Informatics: Artificial Intelligence and Applications*, 2004
- [WEI04] Weippl E., Schatten A., Karim S., Tjoa A. “SemanticLIFE Collaboration: Security Requirements and solutions – security aspects of semantic knowledge management”, in *Proceedings of Practical Aspects of Knowledge Management (PAKM 2004)*, Vienna Austria, LNAI 3336 Springer, pp. 365-377, 2004
- [XACML] “Extensible access control markup language specification 2.0”, OASIS Standard, (available at <http://www.oasis-open.org>), 2005
- [BEL04a] Belsis P., Gritzalis S., Malatras A., Skourlas C., Chalaris I., “Enhancing Knowledge Management through the use of GIS and multimedia” in *Proceedings of Practical Aspects of Knowledge Management (PAKM 2004)*, Vienna Austria, LNAI vol. 3336 Springer, pp. 319-329, 2004
- [BEL04b] Belsis P., Gritzalis S., Skourlas C., Drakopoulos I., “Implementing Knowledge Management techniques for security purposes”, in *Proceedings of the 6th International Conference on Enterprise Information Systems*, Porto, Portugal, Proceedings Vol. 2, pp 535-540, 2004
- [BEL05] Belsis P., Malatras A., Grizalis S., Skourlas C., Chalaris I., “Sec-Shield: Security Preserved Distributed Knowledge Management between Autonomous Domains”, in *Proceedings of the 2nd International Conference on Trust and Privacy in Digital Business (Trust Bus 05)*, Copenhagen, Denmark, LNCS Springer, 2005
- [MAL05] Malatras A., Pavlou G, Belsis P., Gritzalis S., Skourlas C., Chalaris I., "Secure and Distributed Knowledge Management in Pervasive Environments", in *Proceedings of the 1st*

IEEE International Conference on Pervasive Services ICPS 2005, V.Kalogeraki (Ed.), July 2005, Santorini, Greece, IEEE Computer Society Press