Interconnecting Autonomous Medical Domains

Security, Interoperability, and Semantic-Driven Perspectives for Electronic Health Records

BY STEFANOS GRITZALIS, PETROS BELSIS, AND SOKRATIS K. KATSIKAS

The emergence of networked infrastructures and electronic health records (EHRs) has brought new challenges in the field of medical informatics. The potential benefits that can be gained from the realization of interconnected medical domains are undisputed. Timely delivery, where needed, of clinical information can improve the quality of clinical care provision. On the other hand, the collection, storage, and dissemination of personal patient data introduce several major challenges. Healthcare information needs to be accessible by authorized users only, while its fundamental security properties (namely, integrity, availability, and confidentiality) must be retained. Medical information is used primarily for treatment purposes; other uses of pieces of medical information, subject to governance and ethical regulations, include education, research, planning of social services, and public health. In this article we discuss the technological challenges from a security and interoperability perspective toward enabling the interconnection of different medical domains.

Toward Integration of Medical Domain Infrastructures

Mobility of patients, especially those who suffer from chronic diseases, leads to a fragmented EHR in different and distant locations, stored in different types of information systems, with incompatible technologies. The result is the evolution of diverse and heterogeneous versions of a patient's medical record. Enabling access to health records from different locations leads to the provision of faster and better management of patients; additionally, it enables the simplification of procedures from the hospital's perspective. If, for example, the management framework includes storage and retrieval of image files, (i.e. X ray, nuclear medicine, and ultrasound), then the patient can avoid undertaking the same tests twice, while the doctors can formulate an opinion about the patient's condition in less time.

Many countries worldwide are trying to realize distributed health information system infrastructures that enable retrieval from and secure dissemination of EHRs to different locations. Unfortunately, nonuniformity among the structures of medical records in conjunction with the inconsistent regulatory and legal frameworks pertinent to patient data confidentiality that hold in different countries reduces the prospect of fast progress.

Digital Object Identifier 10.1109/EMB.2007.901783

Lately, many European research projects have strived to define distributed architectures that will allow the interconnection of different medical domains [2], [7]. The overall target is to enable transparent retrieval and dissemination of medical information in a secure and efficient manner. For this purpose, several architectural requirements have to be met. These requirements mainly refer to access control models, access control enforcement, interoperability, and use of digital signatures for both confidentiality and integrity of medical records.

Access Control Models

In large organizations data are not owned by individual users but by the organization itself; hence, access to data should depend on the position held by a user in the organizational hierarchy. This is the main justification for establishing a model where access controls are determined for groups of users rather than for individuals. Doing so enables scalability, since every organization owns a very large number of assets, each associated with many pieces of information that need to be accessed by different individual members of staff, with different levels of permissions. Group-oriented rather than individual-oriented access control makes the process more easily manageable.

One of the most widely used access control models nowadays is the role based access control (RBAC) model [9]. The main concepts of RBAC are users, roles, privileges, and sessions. A user represents a human entity or an autonomous agent. A role is associated with the execution of a specific task, while a collection of permissions is assigned to each role, enabling the fulfillment of the obligations associated with such a task. The main RBAC principle is related to the fact that usually users with similar roles need to be accredited for the same actions and need to have the same access rights. By identifying different roles and subsequently relating individuals to roles, the security problem can be significantly simplified. RBAC has become the dominant security model due to its flexibility, its scalability, and its capability to reflect organizational hierarchy; moreover, its parameters can be easily codified. As a consequence of its wide acceptance, RBAC has become a de facto standard.

For distributed authentication, several modifications to the standardized RBAC need to be considered. Such extensions include context-related predicates in role definition (for

example, the domain where a specific user belongs to, parameters related to personal data, etc.). Other extensions also enable determination of time-based restrictions (how long a role is allowed to be activated by a user). 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 security measures.

Policy-Based Management

By "policy," specifically access control policy, we refer to the set of high-level rules according to which we regulate access control. The main reasons for using a security policy are scalability and flexibility from a management perspective. Scalability is achieved by applying the same rules for a large number of assets and flexibility by distinguishing the management procedures from the actual implementation of the access control system; thus, policy may be changed without making it necessary to modify the implementation or to interrupt the operation of the system. A typical enterprise system may include a large number of heterogeneous assets and may also offer services to a large number of users with diverse access rights over these resources. Thus, management also needs to be dynamic and flexible to deal with the evolution of the systems being managed [8]. Medical information systems are exposed to many threats, and are often characterized as sensitive systems, due to the nature of the data they handle. Therefore, it is necessary to deploy a combination of solutions that enables flexible management while preserving privacy.

In order to handle interoperability issues, policies should be formulated and encoded in such a form that their correct application and interpretation are guaranteed. Therefore, policies have to be standardized regarding syntax, semantics, vocabulary, and operation. This can be achieved through the specification of special-purpose schemata, such as Extensible Markup Language (XML) schemata [5]. Among several existing languages for policy specification and access control enforcement we have chosen XACML [1] for the deployment of an experimental infrastructure that allows the interconnection between autonomous domains. Some of the reasons for choosing this language include:

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

The basic modules of the distributed access control framework are (Figure 1):

- the policy enforcement point (PEP), which grants access to roles after receiving an appropriate message from the policy decision point (PDP)
- the PDP, which reasons over a specific access request after evaluating both the requestor's credentials and the request according to the policy in force
- the policy information point (PIP), which is responsible for retrieving environmental attributes, and
- the context handler, responsible for appropriate structuring of messages according to the language syntax and for evaluating context-related parameters from the PIP during the authorization evaluation process.

In more detail, the access control framework operates as follows (Figure 1). The policy administrator is responsible for editing the policy and makes it available to the PDP. When a request for a resource coming from within the same domain appears (Figure 1, action 2a), it is directed to the PEP, which asks the PDP to validate the request for its consistency with the local security policy prior to its execution. The PEP's request to the PDP is constructed as an appropriate XML message and directed to the PDP. Prior to the validation of the request, the context handler and the PIP send additional subject, resource, action, and environment attributes to the PDP. Finally, the initial access request is validated by the PDP and a response message is sent to the PEP, which handles the details for providing authorization to the requester.

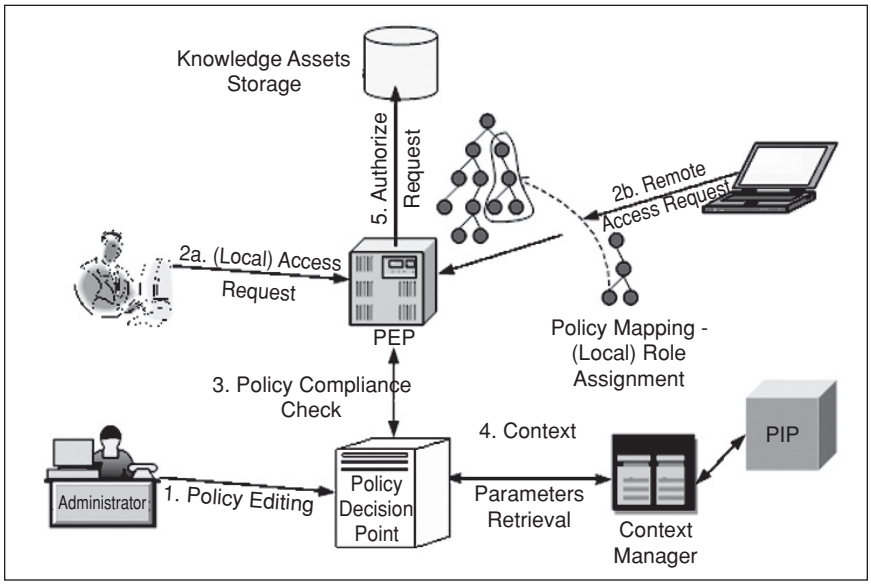

Fig. 1. Authorization process in steps. The policy is being edited by the administrator and stored at the PDP. Local requests are processed by the PEP and PDP and access is granted or denied depending on if the requester should be authorized according to the policy specification. In case of a remote request, an appropriate mapping has to be retrieved first from the special-purpose registry and then the process works as in the case of a local request. The role of the context manager is to retrieve domain-specific attributes and facilitate authorizations for all members of a domain.

Multidomain Access Control Enforcement

Our work extends this single-domain authorization framework to provide support for role and privilege assignment for users belonging to remote domains. This is necessary when users (for example, doctors) need to be assigned privileges to access medical data in other medical domains to which they do not have routine access. In order to achieve this interconnection between different medical domains, several issues need to be taken under consideration.

- Access to data should be regulated by specific generic guidelines, applicable for all the cooperating medical environments.
- While the data access guidelines should be uniform, enforcement points should be autonomous and should have a large degree of freedom in managing their IT infrastructure.
- Dynamic nature of healthcare units coalition. The number of healthcare units that participate in the cooperating schema is not stable. Units can join or leave at any time, increasing the complexity of the overall management.
- Absence of a centralized authorization architecture. Security policies can be defined locally, without the necessity for central management, which would endanger the system's performance by introducing a single point of failure. It would also not be consistent with the distributed nature of the system.
- Transparency to the users. The procedures for accessing medical data, whether these are stored locally or in a remote domain, should be identical for the user.

Every solution attempting to enable intradomain communication should also take into account the demand for scalability support as well as for minimum additional overhead, in terms of computational resources and available bandwidth, on the overall architecture. In order to enforce cooperation between different access policies, we define a policy mapping process that enables roles from one domain to be mapped to another domain. A one-to-one role mapping would introduce a considerable amount of complexity to the overall system, and it would make it difficult to scale. Instead of this, we adopted the following approach [3]: Each organizational hierarchy can be represented as a higher-level policy and codified in appropriate policy language format. For interoperability reasons, as explained in previous paragraphs, we have chosen an XML-based policy representation. In order to capture role relationships present in the RBAC model, we have chosen an XML-based but more expressive than typical XML for policy representation; our implementation thus, utilizes the resource description framework (RDF) [6] language. Now, the problem of mapping roles between different domains is transformed to the problem of mapping different hierarchies, or in other words, different ontologies. Therefore, we introduce a general role hierarchy, with generic roles, to which the local roles of the participating domains have to be mapped (Figure 2). The mapping of each role of a domain to the generic authorization hierarchy can be performed by the coalition administrators, who are aware of all the consequences of an incorrect mapping and of potential information exposure to nonauthorized personnel. In addition, the coalition administrators are responsible for establishing new mappings when necessary. Each domain maintains its policy and the policy mappings toward the global

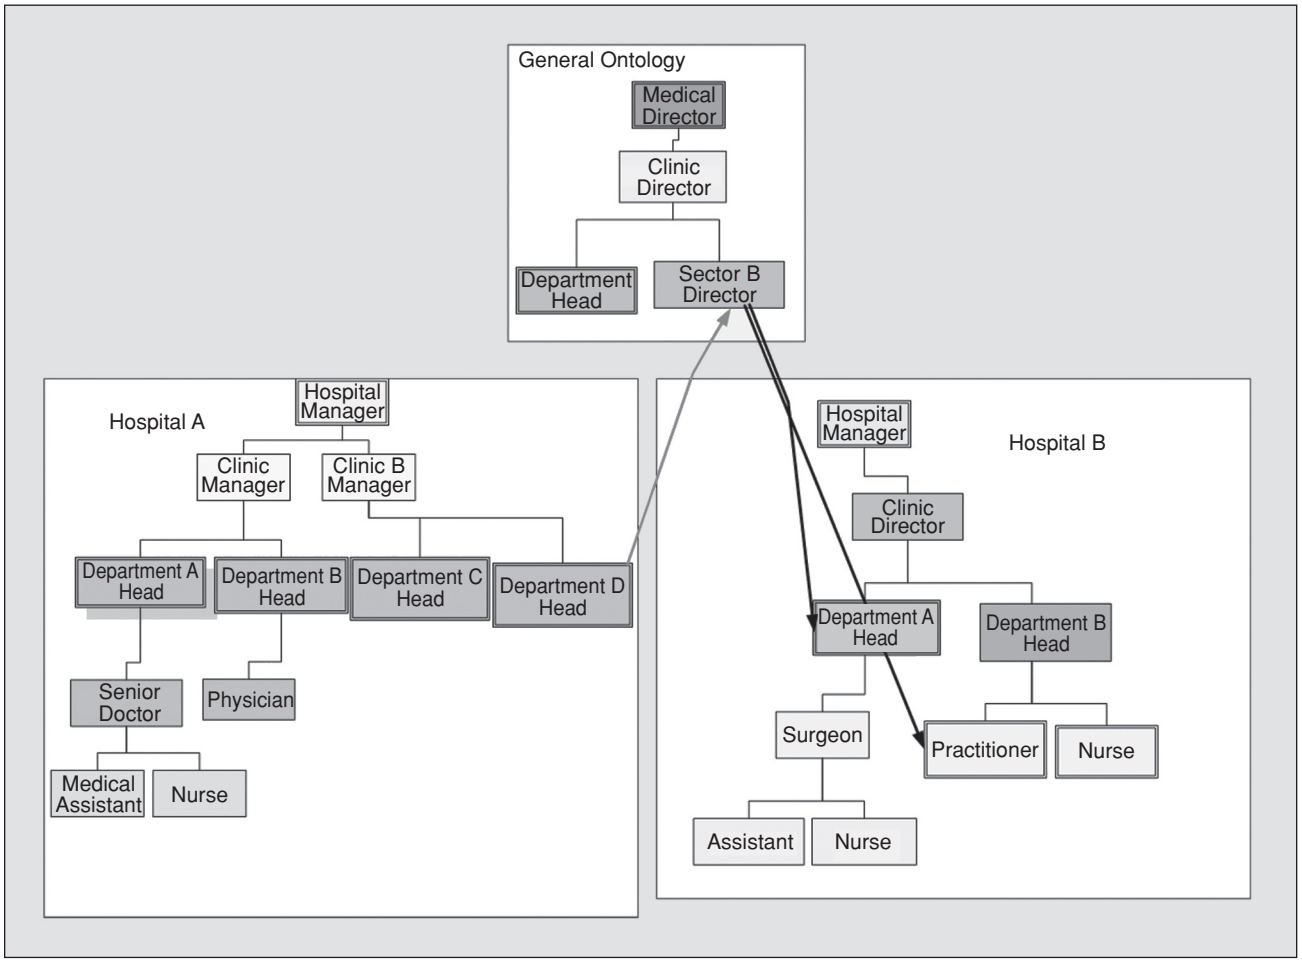

Fig. 2. Mapping roles between different domains. A role from hospital A is appointed through an intermediate mapping to the generic role hierarchy scheme (which acts as a standard for all domains) to a correspondent role belonging to hospital B. Therefore, each user acquires the permissions of the correspondent role, without additional overhead for the system.

authorization scheme in a registry maintained for this specific purpose. The domain's policy, the policy mappings, and the global hierarchy structure can all be replicated to several nodes on the network to avoid the existence of a single point of failure. The global scheme role-hierarchy assumption, to which local policies map, is realistic, since many hospitals define similar but not identical roles. For example, roles defined for hospitals could be Hospital Manager, Ward Manager, Specialized Doctor, Nurse, and so on. Now doctors who work as general practitioners in one clinic could be appointed to work also in another clinic; thus, they would be able to access the same information and have the same privileges. Therefore, a mapping between similar roles could be established in terms of a common role hierarchy, or more precisely in our case, based on a generic (global) ontology role representation.

The authorization process, as described previously, needs to be slightly modified to include the case when a remote access request is directed to the PEP of another domain (Figure 1, action 2b). The request from the remote domain is accompanied by the appropriate mapping to the global role-hierarchy so that the remote PDP can reason according to the prespecified mapping and assign a local role to the remote user. In order to prevent misuse and malicious behaviour by external users, the exchanged messages can be encrypted and digitally signed so that their confidentiality and integrity are preserved and their origin can be easily verified. The deployment of encryption mechanisms for exchanging messages, before (during the composition of the request message) and after authorization has been granted, is necessary not only to prevent unauthorized disclosure or modification of medical data but to allow for the use of strong authentication procedures as well.

Interoperability Issues in Multidomain Environments

The challenge for improvement of the quality and the efficiency of healthcare systems causes their switch to distributed and cooperating structures that can realize the shared medical-data paradigm. This paradigm has to be supported by information systems, which must also consist of standardized, open, and interoperable components.

In order to realize their multipurpose use and reuse as well as interoperability at knowledge level, EHRs have to meet special architectural requirements. Toward this end, communication messages as well as EHRs have to at least be encoded in a commonly interpretable way by all participating domains. HL7 is an international standard that aims at enabling communication between applications provided by different vendors, using different platforms.

Communication is achieved through syntactically and semantically standardized messages. Exchanging HL7 messages can be made through cost-effective channels such as the Internet. However, the use of such communication techniques introduces several risks, such as:

- disclosure of medical data to nonauthorized parties
- unauthorized modification of the contents of the message
- nonrepudiation of the messages is not provided, allowing one of the communicating parties to deny having sent or received a particular message.

Existing cryptographic techniques are powerful enough to provide security services of high quality. Several Internet standards describe techniques that are able to thwart the aforementioned threats and to ensure the integrity, authenticity, confidentiality, and nonrepudiation of messages.

Symmetric encryption techniques, which are based on a common key known to the communicating parties, provide solutions to the problem of confidentiality. However, symmetric encryption presupposes the existence of mutual trust between the communicating parties. Moreover, there is no way to verify the identity of the originator of a particular message among those parties that share the same key. Since a number of users all share the same key, the key's secrecy is difficult to be assured for long periods. As the disclosure of the key of one participant results in total loss of security, it is difficult to ensure security through shared keys for more than two communicating parties, while the key delivery process has to take place through other secure channels.

Because of these inadequacies of symmetric encryption techniques, asymmetric encryption is mostly deployed to encrypt messages in distributed environments. Asymmetric encryption works with a pair of keys, which identify an entity in a unique way. One key (the public key) is known by everyone, while the other key (the private key) is kept secret to the owner. Public key encryption can be used to provide the following security services.

- Confidentiality. Data can be sent from the sender to the recipient encrypted by the recipient's public key. This allows only the recipient to decrypt the message. This will also ensure the integrity of the message.
- Authentication. Data are digitally signed by the sender's private key. The originator of the data can be verified by anyone, using the signer's public key.
- Nonrepudiation of origin. As a consequence of authentication, the signer cannot deny being the originator of a message, because no one else could have signed it with the signer's secret key.

Therefore, the use of HL7 messages, together with asymmetric encryption techniques, seems ideal for achieving security by means of retaining confidentiality, integrity, and nonrepudiation.

Based on the experience gained by several European projects [2], [7], most medical distributed architectures use strong cryptographic algorithms and smart cards as secure user authentication tokens. They usually depend on the existence of a trusted third party or certificate service provider that, among other things, vouches for the trustworthiness of all users with regard to their digital identity and credentials.

Secure information transfer over networks is realized via the secure file transfer protocol (SFTP) that has been approved as an HL7/ANSI standard. SFTP can also be used to transfer images. The implementation choice for HL7 is XML [5], which also tends to become the de facto codification and data exchange standard.

Figure 3 presents a prototype distributed, policy-based, HL7-based message exchange architecture. An example scenario of a message exchange between two different domains works as follows: When a request for access is made by a doctor in domain A, it is directed to the domain authorization point (PEP). The PEP sends an appropriate message to the domain policy server (PDP). According to the credentials that the user has provided, the request is evaluated and a role is granted in the local domain. In case the user wants to search for medical data in a remote domain, the appropriate mappings are retrieved from the policy server, and an appropriate message is created and sent to remote domain B. The domain A authorization point creates and sends a request message to the domain B authorization point, containing a description of the requested resources and the specification of the remote role, which is likely to be granted according

to the predefined policy mappings. The message is signed with the sender's private key, so that B can verify the origin of the request. Then, remote authorization point B uses the public key of the sender to decrypt and verify the authenticity and origin of the message. The domain B server reasons over the request after consulting with its policy server and subsequently retrieves the requested information, forms, and signs the HL7 response with its private key. Then, the domain B authorization server creates a session key to encrypt the retrieved information that will be transmitted between the two domains, and signs it with its private key so that the integrity of the message can be verified; in addition, B signs the message with domain A's public key, so that nobody except A can read the message. The use of a session key is preferable, because symmetric cryptography requires less computational and network resources and is robust enough since the shared key will be used only for this session. The domain A server uses its private key to retrieve the session key and then uses this to decrypt the HL7 response. The response is also validated using domain B's public key. All the interdomain communications are encrypted using the secure sockets layer protocol.

Virtual Ontology Networks

Exchanging messages over networks, particularly when these messages need also to be encrypted or verified for their authenticity and origin, is a resource-consuming process in terms of both computer resources and network bandwidth. Hence, its is highly desirable to reduce as much as possible the number of messages that need to be exchanged in order to retrieve the desired information. One way to do this is to increase the efficiency and effectiveness of the process of responding to users' queries for specific medical information. To this end, we have introduced [4] the creation and maintenance of virtual overlay networks corresponding to branches of a predefined ontology that is used for semantically enriching the multidomain medical environment. Each branch of the ontology will have knowledge information from various sources in the medical domain associated with it.

For example, a user may want to search for records concerning data related to hematology. If a clinic does not maintain such data, then querying the network and receiving negative response after some time reduces the system's response times, introduces additional overhead on both the network and the database server, and decreases the overall system's performance. In order to avoid (whenever it is possible) irrelevant queries to a specific domain, we create an appropriate ontology that can be considered as a set of metadata and metaknowledge that provides a set of concepts and terms to describe the information content of the specific domain. The use of ontologies facilitates the development of interoperable systems due to their flexibility and support for semantic conflicts resolution, and they also provide a means for the standardization of terminologies between different domains. With the presence of an appropriate ontology, each domain can be considered as a virtual overlay network (Figure 3). Each request for specific medical information will be directed to an appropriate (in terms of stored medical content) domain. Instead of processing the query between the nodes and servers of the domain, the ontology is first being queried and, if the query is found relevant, the request is forwarded to the authorization point and subsequently

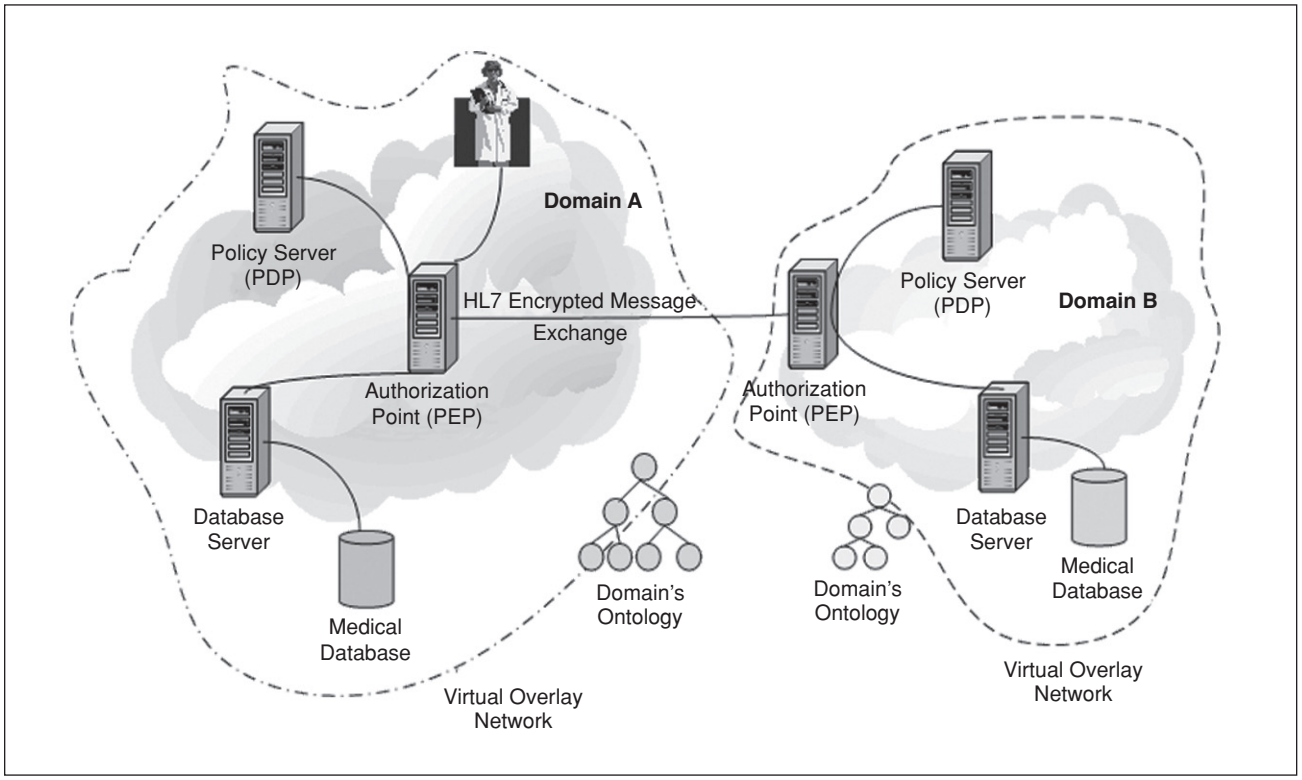

Fig. 3. An example scenario of HL7-based secure message exchange between autonomous medical domains. The presence of virtual ontology networks (VONs) facilitates semantically enhanced querying as well as helps to minimize network resources consumption, since by querying the domain's ontology we are aware if the query is relevant to the domain's assets (for example, we are being informed if the domain stores medical images or not).

to the database servers. This process is particularly important when the medical interconnected infrastructure extends over wireless environments, where device and network capacities are much less than those of wired networks. In these environments the participation of a node to the network affects its power sufficiency and the network topology is subject, to a large extent, to unpredictable variations. Therefore, it is critical to minimize the overall overhead in terms of query processing, while only relevant messages should be forwarded to the authorization points. In these low-resource environments, major categories of the ontology are distributed between nodes forming smaller virtual semantic sets of nodes among the overall set of nodes that participate in the infrastructure.

Implementing interoperable medical information systems is a complicated process because of two fundamental characteristics these systems have: the presence of distributed data sources and their heterogeneity. Information systems heterogeneity may be considered as structural (schematic heterogeneity), semantic (data heterogeneity), and syntactic heterogeneity (database heterogeneity). Syntactic heterogeneity is due to the fact that various database systems use different query languages (SQL, OQL, etc). Structural heterogeneity is due to the fact that different information systems store their data in different structures. Semantic heterogeneity pertains to the content of an information item and its meaning. Semantic conflicts among information systems occur whenever information systems do not use the same interpretation of the data, thus causing serious problems. Ontologies seem a promising concept that can be used for alleviating this problem. In order to use ontologies and technologies available for Web environments, semantic Web languages seem to be the most promising tools; this is why in most cases of building the experimental infrastructure [3], [4], [10], we have chosen the RDF syntax [6] for ontology implementation.

Conclusions and Future Directions

Enabling multidomain infrastructures to communicate and exchange medical information is a very challenging task. Data integration, semantically enriched query formation, resource consumption and network bandwidth, and, of course, security due to the nature and sensitivity of medical data all present major problems. These problems are even more difficult when international data exchange between nations with differing regulatory systems is being considered. International efforts are needed to enable the creation of a widely accepted EHR that will be able to be used between different countries to provide better health services and therefore improve the care of their traveling citizens. Such efforts will complement the technological work that is underway.

Stefanos Gritzalis is an associate professor and head of the Department of Information and Communication Systems Engineering, University of the Aegean, Greece, and director of the Info-Sec-Lab. He holds a B.Sc. in physics, an M.Sc. in electronic automation, and a Ph.D. in informatics all from the University of Athens, Greece. His published scientific work includes several books on information and communication technologies topics and more than 110 journal and national and international conference papers. The focus of these publications is on information and communication systems security.

Petros Belsis is a Ph.D. candidate with the Information and Communication Systems Engineering Department, University of the Aegean, Greece. He holds a Diploma in physics from the University of Athens, Greece; a Diploma in computer science from the Computer Science Department of the Technological Education Institute of Athens, Greece; and an M.Sc. in information systems from the Athens University of Economics and Business, Greece. He has participated in several projects funded by the European Union relative to information systems security, knowledge management, cross lingual content-based information retrieval, etc. His research interests include distributed information systems security, medical informatics, text categorization, and knowledge management.

Sokratis K. Katsikas received the Diploma in electrical engineering from the University of Patras, Greece; the M.Sc. in electrical and computer engineering from the University of Massachusetts at Amherst, United States; and the Ph.D. in computer engineering from the University of Patras, Greece. He now is a professor at the Department of Information and Communication Systems Engineering and rector of the University of the Aegean, Greece. He has authored or coauthored more than 150 books, technical papers, and conference presentations in his areas of research interest, which include information and communication systems security, estimation theory, and adaptive control.

Address for Correspondence: Stefanos Gritzalis, Laboratory of Information and Communication Systems Security, Department of Information and Communication Systems Engineering, University of the Aegean, Karlovassi, Samos, GR-83200, Greece. Tel: +30 22730 82234. Fax: +30 22730 82009. E-mail: sgritz@aegean.gr.

References

[1] Organization for the Advancement of Structured Information Standards (OASIS), "XACML Extensible access control markup language specification 2.0," OASIS standard [online]. Available: http://www.oasis-open.org (accessed Feb. 2005).

[2] B. Blobel, "Authorization and access control for EHR systems," *Int. J. Medical Informatics*, vol. 73, no. 3, pp.251–259, 2004.

[3] P. Belsis, S. Gritzalis, and S. Katsikas, "A scalable security architecture enabling coalition formation between autonomous domains," in *Proc. 5th IEEE Int. Symp. Signal Processing and Information Technology (ISSPIT'05)*, Athens, Greece, Dec. 2005, pp. 560–565.

[4] A. Malatras, G. Pavlou, P. Belsis, S. Gritzalis, C. Skourlas, and I. Chalaris, "Secure and distributed knowledge management in pervasive environments," in *Proc. 1st IEEE Int. Conf. Pervasive Services (ICPS 2005)*, Santorini, Greece, July 2005, pp. 79–87.

[5] T. Bray, J. Paoli, and C. Sperberg-McQueen, "Extensible Markup Language specification (XML)" [Online]. Available: http://www.w3.org/XML/

[6] W3C, "Resource description framework (RDF) specification" [Online]. Available: http://www.w3.org/RDF/

[7] P. Ruotsalainen, "A cross platform model for secure Electronic health record communication," *Int. J. Medical Informatics*, vol. 73, pp. 291–295, 2004.

[8] N. Damianou, N. Dulay, E. Lupu, and M. Sloman, "The Ponder Policy specification language," in *Proc. Policy Workshop 2001*, Bristol, U.K., Jan. 2001, Springer-Verlag LNCS 1955, pp. 18–39.

[9] R. Sandhu, D. Ferraiolo, and R. Kuhn, "The NIST model for role-based access control: Towards a unified standard," in *Proc. 5th ACM Workshop Role-Based Access Control (RBAC'00)*, Berlin, 2000, pp. 47–63.

[10] P. Belsis and S. Gritzalis, "Security control schemes for pervasive medical environments," in *Proc. Workshop Security, Privacy, and Trust in Pervasive and Ubiquitous Computing (SecPerU'05)*, Santorini, Greece, July 2005, pp. 35–43.