W-EHR: A Wireless Distributed Framework for secure dissemination of Electronic Healthcare Records

Petros Belsis, Dimitris Vassis, Stefanos Gritzalis
Dept. of Information and Communication Systems Engineering
University of the Aegean, Karlovassi, Samos, Greece,
{pbelsis, divas, sgritz}@aegean.gr

Christos Skourlas
Dept. of Informatics
Technological Education Institute of Athens, Greece
cskourlas@teiath.gr

***Abstract*—As mobile devices are enhanced continuously with more resources, wireless infrastructures provide support to a growing number of assistive environments. Among the potential domains which can suffice from their deployment, is the e-healthcare sector. The increased sensitivity requirements imposed by the US and EU legislation, urge towards the design and incorporation of strict security standards in the deployment process of wireless e-healthcare infrastructures. Even though mobile devices are characterized by an increase in processing and autonomy capacity, the necessity to encrypt all communications consists of a significant burden. In this paper we present W-EHR, a wireless architecture that enables authorized medical personnel to access medical records in a secure and transparent manner, utilizing an agent based architecture. Ubiquitous access to medical information - within the hospital range - is achieved using mobile devices; security management is achieved using a hierarchical policy based approach in compatibility with the IETF policy-based management model. In order to comply with the strict privacy requirements, all the transactions are encrypted using a hybrid approach that employs symmetric and asymmetric encryption techniques. We describe our experiments that prove the effective operation of our infrastructure (by means of device resources consumption and network bandwidth). In our calculations we consider the encryption overhead when transmitting medical records in a compatible with medical codification standards form.**

Keywords-component; Systems Security; Wireless Security; Electronic Healthcare Records

I. INTRODUCTION

The rapid invasion of wireless devices in many aspects of our lives, has led to the creation of a new computing paradigm. Among else, the e-healthcare sector may benefit from the development of wireless e-healthcare infrastructures. Lately, a lot of research and industry projects aim towards the development of environments that allow dissemination of medical information [1][2][3]. Due to the imposed by legislation strict privacy requirements, the design and implementation of wireless medical environments becomes a complicated task. In this paper, we describe a prototype architecture that allows the provision of accurate - and in secure manner - patient information to authorized personnel within a W-LAN, that spans over the boundaries of a hospital. Thus, doctors acquire patient related information while being close to the patient. Querying for a patient record, authentication and access control enforcement are transparent processes, performed by means of software agent applications, which run at the mobile device. In order to facilitate security management we have adopted a hierarchical policy based management approach that distinguishes two roles between the mobile nodes: Manager Nodes (MN's) which are devices held by doctor's and are responsible to perform more advanced operations (such as to provide access to the data stored in their repository or facilitate authentication for their holder using the agent based application) and Terminal Nodes (TN's) which are allowed to perform simple operations (such as receive information about an emergency situation or acquire access to less critical information) and which are usually held by support medical personnel (nursing personnel). For privacy management, we adopt a hybrid cryptography approach that uses symmetric-key and public-key cryptography techniques; thus, we enable secure transmission of critical information without requiring excessive processing resources in respect to the mobile device's limited capabilities.

The remainder of the paper is organized as follows: related work and background literature is studied in Section 2. Section 3 analyses the basic principles that directed our implementation choices and describes the modular software components of our wireless electronic health records management framework. Section 4 provides technical details in respect to our framework and describes our performance measurements. Section 5 provides concluding remarks and directions for future work.

II. RELATED WORK

Electronic medical record's management attracts increasing interest and sets the scenery for the establishment of ubiquitous acquisition of medical information; as a result, patients and medical personnel may benefit from the improved services of wireless medical infrastructures. A lot of projects have lately focused on providing efficient solutions to various aspects of medical records management [1][2][3].

Wireless mediCenter [3] is a system for management of electronic medical records and delivery through secure LANs

978-1-4244-4530-1/09/$25.00 ©2009 IEEE

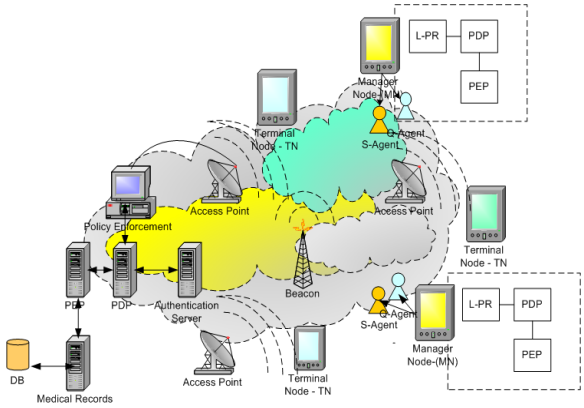

Figure 1 Overall architecture of the distributed W-EHR architecture. The MN nodes contain instances of the policy management module. The beacon transmits signed messages facilitating location based services for the devices.

or high-speed wireless connections. It provides different portals for doctors and patients in order to achieve classification of access permissions. The restriction though to connect through the portal is a serious burden to the user.

PatientService [4] is a trust-based security architecture that enables medical records management in pervasive environments. In this approach access to medical information is provided to a set of users which hold a PDA that keeps the policy in a smart card. In our approach we issue the request form the PDA while the policy evaluation is not performed by the PDA itself. Moreover we attempt to evaluate our approach by performing simulation experiments.

Pervasive healthcare [1] is a project that aims mainly towards context-driven management of medical information using special interfaces and sensors attached next to the patient's bed and provide information to the doctor in respect to the patient's condition. Our aim is to provide information at the authorized medical personnel's mobile device, according to the security classification of the information requested and the role assigned to the medical staff member. We do not assume the existence of special monitors next to the patient's bed, while using an agent based approach we enable information acquisition in a transparent manner.

III. DISTRIBUTED FRAMEWORK ARCHITECTURE

In order to create a distributed architecture that allows secure dissemination of medical records, we adopted a hierarchical, policy based management approach. Policy based management ensures interoperability and scalability features, enabling also uninterrupted system operation during policy execution and while performing policy updates, within the wireless infrastructure. According to our organizational model, which distinguishes different medical personnel roles (with different security permissions) and devices with different processing capabilities, we adopt a hierarchical approach that distinguishes three types of network nodes:

The central nodes (CN's) that store the medical database and are responsible for authentication and access control enforcement and two types of mobile nodes with different processing and access control capabilities: the Manager Nodes (MN's) which are assigned with more advanced tasks and Terminal Nodes (TN's) which lack in respect to MN's in resources and are assigned secondary tasks.

The CN's are responsible for the operation of the policy management module and are responsible to maintain the medical database; they are characterized by adequate processing as well as network bandwidth capabilities. They have also installed the different policy management modules: the Policy Decision Point (PDP) which loads the policies form the Policy Repository and evaluates the request against the recorded policy; the Policy Enforcement Point (PEP) which accepts requests, constructs an XACML [9] compatible request message, sends it to the PDP and finally enforces the decision when it arrives from the PDP. Authentication is performed through an LDAP server which evaluates the medical personnel's credentials (encoded as X.509 certificates) and issues a SAML [10] assertion which can be further used for identification in every future transaction with the access control enforcement module, providing thus a Single Sign-On (SSO) mechanism. In respect to mobile nodes, we distinguish two organizational roles which characterize also the operation of each node: a) Manager Nodes (MN) are devices with more processing capabilities and RAM memory and are held by doctors; b) Terminal Nodes (TN) are devices with less processing power capabilities and are supplied with a lightweight implementation that allows simple operations, like informing a medical assistant about a patient's medication and when it should be scheduled.

The software installed on MN's includes a local PDP and PEP module which allow enforcement of local (as recorded in the device) policies, enabling thus access to other doctor's to the device's local repository. On the contrary, TN's perform only simple operations, such as informing nursing personnel about an emergency or providing details about a patient's pharmaceutical prescription and the time that this medication is scheduled; a TN is never allowed to access sensitive medical data.

Fig. 1 depicts an overview of the overall system's architecture. Each MN contains a copy of the PEP and PDP modules as well as a local policy instance stored at the device's policy repository. A MN also has installed a pair of software agents developed using the JADE agent development platform [5] and the LEAP libraries that allow lightweight implementation for mobile devices. The two agents are responsible for two tasks: a) to query the medical database and b) authenticate and interact with the policy module on the doctor's behalf.

Technically, we employ software agents in order to facilitate interoperation between the different software modules. Agents communicate using the FIPA-ACL Agent Communication Language [6]. ACL (Agent Communication Language), as its name suggests, focuses on the structure and communication related attributes of a message, such as sender/recipient address, message type (e.g. assertion, request, query etc.), ontological commitments, supported content languages and interaction protocols [7]. ACL messages in their payload carry the necessary information, encapsulating thus a query to the database (For the Q-Agent) or an access control request (for the S-Agent). Figure 2 depicts the internal implementation

details of the agent based application and the interoperation with other modules using network and agent protocols. Fig. 2 presents also other platform specific agents such as the DF agent which provides yellow pages services to the S-agent and Q-agent. Another also platform specific agent is the AMS agent which acts as agent coordinator.

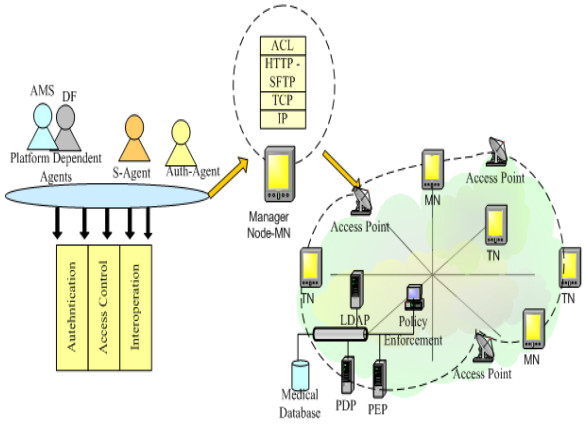

Figure 2 The internal structure of the agent based architecture and its interaction with the other modules within the wireless framework

The Query Agent, in its message carries the request for the medical record of a specific patient; the Security Agent, directs a XACML request that queries the security module for the specific medical record. This message will be directed to the Policy Enforcement Point for further evaluation. The message transfer at the lower level is enabled using TCP/IP and secure protocols such as the SSL to enable encryption for privacy reasons. The S- Agent is responsible for user authentication and access control on the medical personnel's behalf. It retrieves (from the memory-card inserted in the PDA) the doctor's private key and initiates the authentication process by communicating with the authentication server. As aforementioned, in order to keep the resources consumption of the mobile device low, we adopt a hybrid approach that utilizes symmetric and asymmetric cryptography techniques. This is being done since even though most devices support at least 128-bit key length encryption, the use of asymmetric encryption requires excessive processing times and resources. Accordingly, the server issues a SAML assertion which will be used next by the S-Agent for all authentication purposes. The server in addition signs a shared key using the doctor's public key and sends it to the doctor's PDA. This shared key will be used next to encrypt all the necessary information using the SSL protocol.

Both TN and MN nodes are able to identify whether they reside within the clinic or in an unknown environment, with the aid of a beacon (Fig.1) which sends signed messages identifiable by each device when compared to a number of stored signed (within the smart card) messages. Thus, we prevent unauthorized transmission or reception from the device when it resides outside pre-settled space boundaries.

A. Techical Specifications and Performance Evaluation

The distributed policy authorization module is realized by means of object-oriented software architecture, using Java. Communication between the different policy management software modules is achieved using the Java-RMI model. The software agents where implemented using the JADE agent development framework and the LEAP libraries that allow lightweight implementation for mobile devices with limited capabilities.

In order to test the validity of our approach we have also evaluated its performance through simulation in Pamvotis WLAN simulator [8]. We furthermore choose the widely accepted between the medical world HL7 standard to encode the exchanged information. We assume an IEEE 802.11 wireless channel of 1Mb/s, which is capable of covering a range of up to 300m of indoor environments. The IEEE 802.11 protocol is suitable for the transmission of messages originating from HL7 and ACL-based applications, as it supports high data rates and combines the advantages of mobility and packet switching, making it suitable for IP-based mobile devices such as PDAs and 3G/4G mobile phones.

For the Manager Nodes, two basic types of messages exist for each query in the database; ACL messages, containing the link for the transaction of the information, and HL7 messages, containing the real information. We assume that the application packet payload size obeys a uniform distribution with a mean value of 280Bytes. Concerning the packet overhead added from the SSL protocol, this is about 16% of the packet payload [14], meaning 48Bytes. Adding the TCP overhead (32Bytes with timestamps included) and the IP overhead (20Bytes), we have an IEEE 802.11 MAC datagram of 380Bytes. As for the ACL message, we assume that the packet payload size obeys a uniform distribution with a value of 100Bytes. Adding the SSL and TCP/IP overhead, we have a MAC datagram of 168Bytes.

Concerning the traffic generated from Terminal Nodes, we assume simple database transaction messages with a packet payload that obeys an Exponential distribution with a mean value of 200Bytes, which refers to a MAC datagram of 252Bytes by adding the TCP and IP header overheads.

Additionally, we assume that the number of transactions each Manager Node and each Terminal Node perform, obey a Poisson distribution with a mean of one message every two minutes.

Finally, we assume that the network is consisted of 300 Manager Nodes and 300 Terminal Nodes. The simulation results concerning the above configuration are outlined in the rest of this section. Figure 3 depicts the aggregate traffic (system throughput) generated by messages sent from Manager and Terminal Nodes. As we can see, the traffic is minimal compared to the wireless channel capacity. Hence, if other applications run on the same network (e.g. medical image downloading or interactive applications such as VoIP), our proposed architecture does not affect their performance.

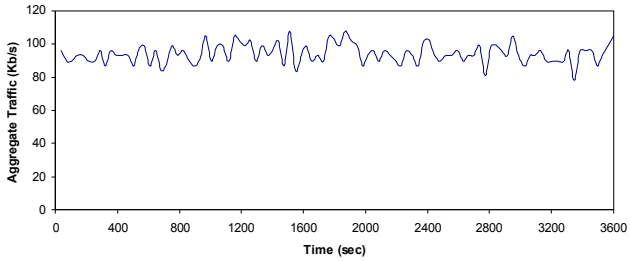

Figure 3 Aggregate Traffic versus simulation time

Figure 4 depicts the service delay for a network consisted of 600 users. What we mean by service delay is the delay from the moment a user sends a query to the database server, until the moment he receives the HL7 message from the database containing the information requested. Note that the processing delay (e.g. delay faced on the database server for performing the database transaction) is not taken into account, so basically the service delay refers to the sum of the packet delay of an ACL message and the packet delay of an HL7 message.

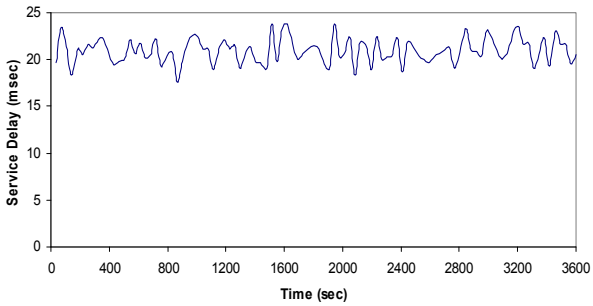

Figure 4. Service delay versus simulation time.

We observe that the service delay is less than a second, even for the best effort class considered in our case. Figure 5 depicts the service delay versus an HL7 message size in ASCII characters. As we can see, the delay is very low. The system gets saturated for messages longer than 16000 characters.

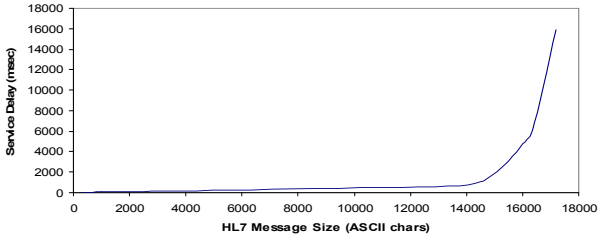

Figure 5. Service delay versus HL7 message size.

Finally, Figure 6 depicts the service delay of Terminal Nodes that refers to the duration from the beginning of a database transaction until the reception of the information from the database server. As we can see, the delay is very low, allowing fast transactions, even if other services run on the network.

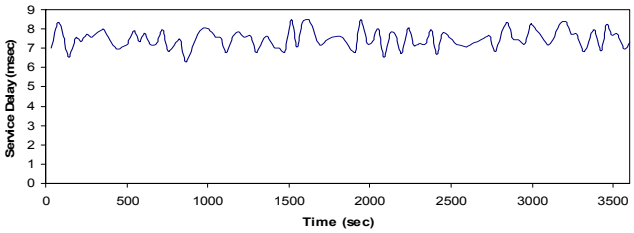

Figure 6 Service delay of Terminal Nodes versus simulation time.

IV. CONCLUSIONS

In this paper we presented W-EHR, a wireless distributed framework that allows secure dissemination of electronic medical records within a wirelessly covered area using mobile devices. For security management we have adopted a hierarchical policy based approach that enables the security management of devices with different roles within the framework while providing our architecture with scalability features. Compliance with privacy regulations has been achieved using a hybrid approach that utilizes symmetric and asymmetric cryptography techniques. In order to further facilitate the medical record retrieval and access control enforcement process, we have created a lightweight agent based application that is installed on the medical personnel's devices. The validity of our approach was tested by performing simulations in which we have estimated the number of queries and the resource consumption while considering the cryptography and HL7 overhead.

We intend to apply our scenario to a wider extent and to measure our platform's performance through extensive testing and recording of efficiency parameters for a large number of users and devices.

REFERENCES

[1] http://www.pervasivehealthcare.dk/projects/index.html

[2] http://www.eecs.harvard.edu/~mdw/proj/codeblue/

[3] WirelessMedicenter. http://www.wirelessmedicenter.com/mc/glance.cfm

[4] A. Choudhri, L. Kagal, A. Joshi, T. Finin, and Y. Yesha, "PatientService: Electronic Patient Record Redaction and Delivery in Pervasive Environments", Proceedings of the Healthcom 2003 Fifth International Workshop on Enterprise Networking and Computing in Healthcare Industry, Santa Monica, USA, June 2003

[5] JADE Software Agent Development Platform, http://jade.tilab.com/

[6] Foundation for Intelligent Physical Agents. FIPA Speci-fications, http://www.fipa.org/specifications/index.html, 2005.

[7] V. Zafeiris, C. Doulkeridis, P. Belsis, I. Chalaris "Agent-mediated Knowledge Management in Multiple Autonomous Domains", Workshop on Agent Mediated Knowledge Management, Univ. of Ulterecht The Netherlands, July 2005, pp. 97-112

[8] *The Pamvotis WLAN Simulator, I*nformation available online at www.pamvotis.org.

[9] Moses et al, eXtensible Access Control Markup Language specification, v.2 Technical Overview, May 2004. available: XACML Oasis TC Homepage, http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml

[10] Hughes et al., *Technical Overview of the OASIS Security Assertion Markup Language (SAML) V1.1.* OASIS, May 2004 http://xml.coverpages.org/saml.html