Security and trust in virtual healthcare communities

Anargyros Chryssanthou
Greek Data Protection Authority
Auditors Department, ICT Auditor
1-3, Kifissias Avenue, Ampelokipoi
+30 210 6475677
achrysanthou@dpa.gr

Iraklis Varlamis
Harokopio University of Athens
Dept. of Informatics & Telematics
89, Harokopou St
+30 210 9549295
varlamis@hua.gr

Charikleia Latsiou
Greek Data Protection Authority
Auditors Department, Lawyer
1-3, Kifissias Avenue, Ampelokipoi
+30 210 6475698
clatsiou@dpa.gr

ABSTRACT

The main purpose of a virtual healthcare community is to enforce members' collaboration and allow them collectively conduct health care activities. Patient monitoring and medical consultation and support are the most popular activities inside health care communities. They bring together medical experts and patients and require confidentiality, reliability and trust in order to be successful. An examination of existing virtual communities for healthcare leads to the conclusion that many of them fail to meet requirements for building trust. Several ethical, legal and technical issues must be considered in order to build a trustful community. This work presents the architecture of a virtual healthcare community portal with emphasis on the issues that help building trust inside the community. With a set of hypothetical usage scenarios that challenge trust in the community we uncover healthcare community's pitfalls and illustrate the solutions provided by the proposed architecture.

Categories and Subject Descriptors

K.6.5 [**Management of computing and information systems**]: Security and Protection – *authentication, insurance, unauthorized access*

General Terms

Design, Security, Human Factors, Legal Aspects.

Keywords

Security, Trust, Virtual Communities, Health Care, Tele-consultation

1. INTRODUCTION

The progress in telecommunication technologies has removed several time and distance barriers and allowed virtual communities to flourish. In virtual healthcare communities patients submit online requests for advice and share their problems and knowledge, doctors cooperate with each other, supervise and support their patients. Doctors, patients and care givers are all members of the same community and collaborate in order to virtually manage the illnesses and improve the quality of patients' life. Specialized healthcare communities, such as self supportive patient communities that promote peer to peer patient communication and medical research communities that support the collaboration of medical professionals can be facilitated by the virtual organizational model.

The most significant issue in healthcare applications is the protection of patient's medical data from unauthorized access. In pervasive and web based healthcare applications, medical data is transferred via wireless networks and/or across the web, so specific attention should be drawn towards building and meeting the appropriate security requirements. Therefore, it is important to protect the confidentiality of sensitive medical data, maintain its integrity and ensure that sensitive medical data is always available to the rightful data holder (patient or doctor).

Trust is another important issue in healthcare communities and requires more than a secure technological solution. Members of a virtual medical community need to trust each other and to be confident for the secure, reliable and lawful operation of the community. As described in [4], building of trust is a continuous process that comprises several repeating steps: achieving an appropriate security level for medical data in terms of authentication and user's certification, defining a strict user policy with roles, access rights and limitations among community members, providing a flexible identification mechanism, which preserves anonymity whilst guaranteeing identity truthfulness. Additionally, in patient monitoring cases, the community must respond quickly and reliably upon emergencies.

Continuity is the last but most important issue for any virtual community. System maintenance, based on auditing and vulnerability testing is necessary for the stability of the community infrastructure, from a technical point of view. A reputation system may help to elicit good behavior, encourage knowledge sharing among individuals and strengthen members' bonds to the community.

This work presents the desired architecture for a virtual healthcare community portal. It enlists potential hazards for the members of the community and provides suggestions on the technological infrastructure, the operation policies and additional certification mechanisms and other functionalities that increase trust. The work capitalizes in the technological and social dimensions of trust and describes a holistic approach in building and maintaining a trustful and secure virtual community for healthcare.

PETRA'09, June 09-13, 2009, Corfu, Greece.
Copyright 2009 ACM ISBN 978-1-60558-409-6...$5.00

The following section refers to related works that partially cover the community needs for security and trust. Section 3 illustrates the desired structure of the virtual community platform. Section 4 enlists all potential threats for the community from technical, ethical and legal aspect and section 5 presents several community threats, focusing on the mechanisms that confront them. Finally, section 6 presents our conclusions from this work.

2. RELATED WORK

Several projects concerning the development of virtual healthcare communities that support the pervasive participation of patients (e.g. through wireless monitoring devices) have attracted national or private funding. CIGNA has already announced a virtual community for nutrition and healthcare, which is situated on a Second Life island [16]. The EU-funded project Saphire [13] has integrated wireless medical sensor data with hospital decision support systems in an attempt to provide remote monitoring of patients at their homes. Several more research works on healthcare delivery [6], patient peer-support [21], virtual disease management or medical research and collaboration through virtual medical communities, have been found in the literature.

The quality of patient services is strongly related to the availability and quantity of medical information. In order to increase the quantity of medical information without burdening the patient, several sensors based monitoring systems have been designed that allow continuous recording of patients' status, such as CodeBlue [14], Scalable Medical Alert Response Technology (SMART) [7], MobiHealth [11] etc. In order to increase the availability of medical information, patients should be persuaded to trust the monitoring and information recording infrastructure, their doctors and the health-related community as a whole.

The first step in this direction is to allow members to secure their data and selectively provide access to them. In the Cassandra trust management system for medical communities [3], access control is based on the member's role in the community. However, each data owner is able to define the access rights on her personal data using the prototype role-based access (RBAC) model. Access rights are validated using a Datalog extension with constraints. XML-based models have also been employed for the same task. XrML [26] allows the definition of rights and granting policies with validity restrictions. XACML [19] is another model for defining conditional access and deny policies and policy combination rules for resolving conflicting policies (e.g. First-Applicable, Deny-Override, Permit-Override). XACML does not support delegation and is thus not well suited for decentralized authorization. Finally, the Security Policy Assertion Language (SecPAL) [2] is another XML-based model, which builds on the notion of tunable expressiveness introduced in Cassandra.

Solving the security issues that relate to the wireless or wired transmission of data [18] and the legal and ethical issues concerning the confidentiality of patient data [23] are not always adequate for building trust in the healthcare community. According to [12], trust is subjective, bi-directional though asymmetric, non-transitive, context dependent, dynamic and time-dependent. A trust management mechanism that keeps record of the members' reputation inside the community and continuously updates it by analyzing other members' feedback can be useful in this direction.

3. A VIRTUAL COMMUNITY FOR MONITORING AND TELE-HEALTHCARE

Virtual communities refer to groups of people that collaborate, discuss their issues, share experiences, consult with experts, provide and request for support using telecommunication technologies. Virtual healthcare communities, employ advanced and pervasive ICT technologies in order to offer ubiquitous medical services to their members. Elder members, home care patients or members with chronic conditions, utilize different types of health care services at different points in time, bridging geographic distance and time constraints [4].

3.1 Community members

The active members of a virtual community in health care comprise patients, doctors, as well as people with interest in the community issues, such as patients' family members, researchers etc. [25] Members have different roles depending on their needs and expertise: patients and family members undertake facilitator roles, healthcare professionals become moderators for discussion and contents, facilitators and mentors for the community members. The system administration of the community is usually performed by IT experts who must be trustful community members.

In complement to the community members, several people, in the community background, guarantee the smooth operation of the community and the uninterrupted delivery of services. The IT staff that technically supports the community, the employees of the telecommunication services provider and the directors of the organization, company or hospital that hosts the virtual community are persons that do not actually participate in the community but play a key role in its secure operation.

3.2 Community activities

An overview of the interactions inside the healthcare community is presented in Figure 1. Health status signals are collected and transferred to the community portal, while patient members request for advice, diagnosis or treatment suggestion etc. The doctor from inside the hospital is able to access the patient's record and make a diagnosis. The doctor can also consult the patient directly, based only on the patient's medical status signals. The hospital keeps record of patients' profiles and history, doctors' diagnoses, and of all requests and advices exchanged in the portal.

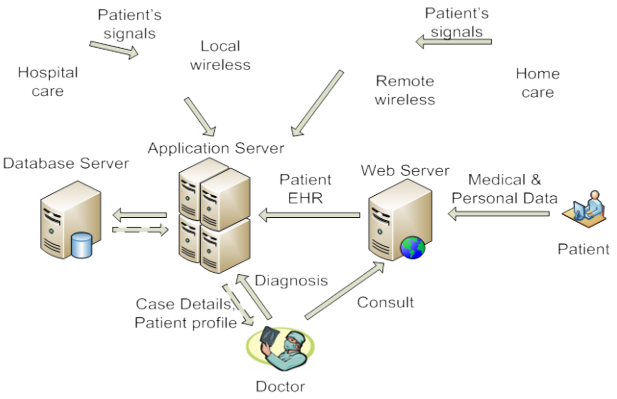

Figure 1. Overview of the community interactions

3.3 Multi-tiered security model

The process of achieving an adequate level of security in the networked environment of the virtual community is twofold. First, the ***internal layer*** of the community needs to be secured. This mainly comprises the database server, where the sensitive medical data reside and the ***application layer*** where the user requests are served. Second, we must certify that users in the ***community perimeter*** (mainly patients) have an adequate level of security.

The intranet/internet model used in the past was based on the notion that a firewall is adequate to secure the inside perimeter of the system (intranet). However, the advent of pervasive and ubiquitous computing created new challenges for computer security professionals. People are connected to the internet from anywhere; make use of advanced community services without being aware of how to interact with them. Third party applications, devices and networks interact with the same services and access the same resources. As a consequence, an interface to these applications is necessary and a ***presentation layer*** for interacting with individuals is required.

The iMedik telemedicine system [15], presented a four-tier architecture comprising: a web server on the demilitarized zone (DMZ), a web proxy layer in front of the firewall and finally the application server and the database protected behind the firewall. The first level of authentication is performed on the web proxy layer. This four-tier model can solve authentication and security issues of the medical community.

In the proposed structure, the proxy layer checks the validity of a user session (whether the user is authenticated or not), the presentation layer (web server) validates the user input and the application server checks whether user's permission on the requested data is sufficient. The user is authenticated outside the perimeter and any invalid attempts will fail grace to the firewall (Fig.2, 3). Moreover, typical web attacks such as cross-site scripting or sql injection, which may be performed by a valid community member, can be detected at the presentation layer. Finally, user permissions and access rights on medical data can be verified at the application server level. This multi-layer approach keeps unauthorized users outside the community's perimeter and guarantees that authorized users cannot gain invalid access to medical profiles or access the database in a disallowed mode.

Figure 2 presents the overview of the desired community architecture, with all the suggested servers and other security mechanisms. The subsections that follow, explain the details of this architecture and the reasons behind each decision.

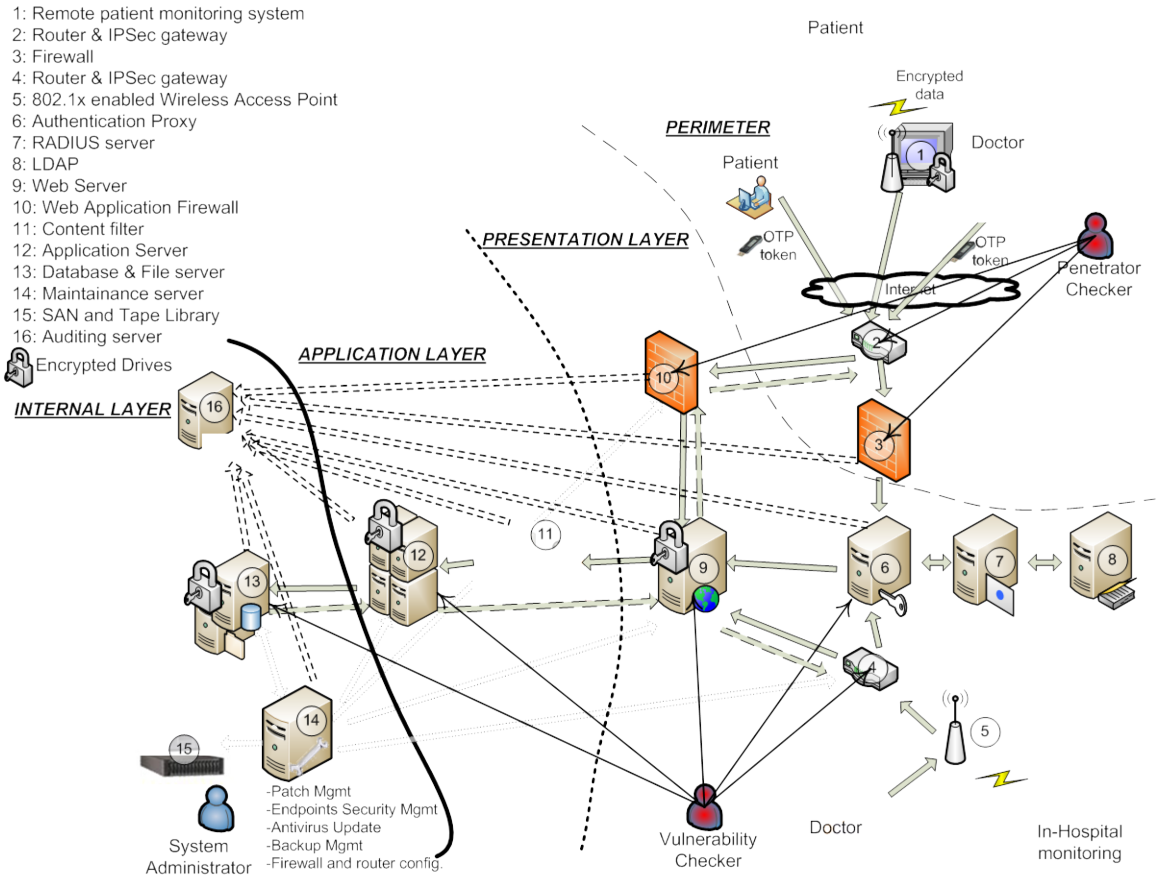

Figure 2. Security enabled architecture

3.3.1 *Community perimeter*

The identification of the virtual community members can be performed outside the web server by means of an e-token that connects to a Remote Authentication Dial In User Service (RADIUS) server (Fig.2, 7). The RADIUS server can be part of the community infrastructure or (preferably) belong to a certification authority such as Verisign. This stateless server integrates with a local directory (Fig.2, 8), where the identity store of the virtual community will be kept. An authentication proxy (Fig.2, 6) in front of the RADIUS server will be the enterprise validation engine. RADIUS proxy-based deployment is ideal for network applications such as VPN remote access. An e-token device will allow single sign-on to community services, since the users' id will be kept into the identity store and in combination with the OTP, generated by the e-token device, will be used for user authentication in every subsequent login.

Audit logs of each transaction will be kept for non-repudiation purposes and for incident handling purposes in the auditing server (Fig.2, 16). Namely, each and every transaction can be traced back to the person responsible, so there will be no disputes in case of security incidents concerning medical data.

Doctors from inside the hospital and patients, who are treated inside the hospital, are able to enter the community through an 802.1x enabled wireless access point (Fig.2, 5). Doctors' mobile devices and the sensors that monitor patient activity are connected to a wireless base station configured to use 802.1x protocol and all traffic is forwarded to the authentication proxy. The proxy is configured to require 802.1x from all clients connecting through the particular wireless router and ignores any other incoming connection. The identity of the client device is forwarded to the Radius Server (authentication server). The authentication server sets up an EAP-TLS session with the client using digital certificates for mutual identification. If valid digital certificates are used the client is successfully authenticated.

In the case of home-care applications, remote monitoring systems (Fig.2, 1), consisting of wireless sensors attached in the patient's body need to securely communicate with the application server of the community. Off the shelf wireless sensor platforms with security features, such as TinyOS can be employed given that they offer software or hardware encryption of wireless transmitted signals [27]. The encrypted signals, which are collected to the base of the remote monitoring system, can be forwarded via the secure router (Fig.2, 2) to the application server. Devices that do not support data encryption must be cable-connected to the base.

3.3.2 Presentation layer

The presentation layer is the target of multiple forms of web attacks. Malicious users attack web applications using cross site scripting, SQL injection, HTTP request smuggling, etc. In the proposed architecture, a filtering module (Fig. 2, 11) will "wash-out" malicious user inputs and will block several of the aforementioned attacks. The filtering module processes user input and ensures that user requests through the web server (Fig. 2, 9) do not attack the application server (Fig.2, 11). Commercial web application firewalls (Fig.2, 10) can be employed to perform input filtering, record all traffic that is directed to the database server and distinguish between legitimate requests and potential attacks. The system administrators can configure routers and firewalls (Fig.2, dotted arrows) using signatures of known attacks, provided by the firewall's manufacturer, and create additional signatures for illegitimate traffic.

3.3.3 Application layer

Access to the application server is allowed only to authenticated community members, third party applications and devices. The primary aim of the security mechanisms in this layer is to guarantee that users are properly identified and have access only to the data they are allowed to. The role of the application server (Fig.2, 12), in terms of security and access control, is to verify that authenticated users are authorized to access the requested data.

The ***implementation of access policies***, the ***definition of community roles***, and the ***clarification of access rights and restrictions*** for each role are required in this layer. Concerning the access control, access credentials and initial roles can be distributed by the authorities of the community. The configuration of the application server is performed in the maintenance server (Fig.2, 14) by the community administrator. Then patients can define which doctors can access their private data or which members can take part in a private held conversation using a semantic role-based policy. This access policy definition will be required by each upon registration.

All the above, can be made possible by adopting a flexible access policy model. The simplicity and extensibility of SecPAL [2] along with its PKI-based, SOAP encoded infrastructure for exchanging policy assertions, renders it ideal for access policy management in the distributed virtual community environment.

For example, when a patient needs to undergo a surgical operation, he can define an access policy that asserts full access to the entire EHR file for his doctor and only view permissions for the nurses. In the same policy model, the patient asserts that the doctor has the right to grant the full access privileges on his EHR file to the members of her medical team so that they can assist her during surgery. If the doctor decides to grant full access to an assistant or to any member of the community, all the respective asserts are activated to decide whether granting of privileges is allowed. The SecPAL model allows to define the duration of assigned roles (e.g roles assigned to the members of the surgical team are valid only during surgery and recovery). Upon recovery of the patient the access rights are automatically revoked.

An overall access and behavior policy is supplementary to the access control mechanisms mentioned earlier. The implemented access control model will be explained to the user upon registration and will be available as a written and electronic document. This written document will explain everything, from login procedures, password quality, privacy rights to user roles and credentials. In this way, the users have no excuse for violating other user's rights or the policy in general and the community is protected against users' misbehavior.

Security applies to any type of data, whether sensitive or not. In the case of sensitive medical data the required level of protection is even higher. For example, processing of sensitive data in Greece requires, apart from the users' permission, an extra permission from the Greek Data Protection Authority (Greek Law, 2472/97, article 7, par. 2). An access policy, a security plan, a code of conduct and a risk analysis document are required by the Greek Data Protection Authority in order to grant permission for processing sensitive data.

3.3.4 Community internal network

In order to improve community trust, we must first define members' responsibilities and consequently certify that members carry them out properly. Auditing can assist in this direction. Every single user action, either local or remote, must be logged (Fig.2, dashed arrows). Logs must be securely stored, in an encrypted format, for a period of time, which will be defined after a proper risk assessment on the system. Access to the logs will be allowed only in presence of at least two administrative entities of the community with different roles (for example an IT administrator and a hospital manager), in order to avoid “accidental” or "intentional" data loss. Administrator’s access must also be logged, in order to avoid abuse of administrative access. Logging of all database transactions will allow back-tracing to the perpetrator in case of any improper data access.

All servers must run antivirus software, which is frequently updated. Community members must also have constantly updated antivirus software on their systems, in order to secure both ends of the communication. It is advisable that the server systems are equipped with encrypted hard drives, which can be read in presence of hardware tokens. Endpoint security must be in place so that no external devices (e.g. USB drives) can be connected to the sensitive modules of the system.

Finally, a disaster recovery plan must be designed in order to ensure that in case of a disaster, the virtual community will be operational shortly.

3.4 Security maintenance processes

As a secure infrastructure is important for the operation of the virtual healthcare community, we should make sure that the appropriate security level is attained at all times. Periodic checks are expected to detect new security vulnerabilities and confront evolving attack techniques. More specifically, the security of the authentication mechanism should be checked, the effectiveness of the application firewall must be validated, the security of the authorization process must be checked, auditing mechanisms must function properly, etc. Furthermore, patch management must be applied, so that all servers and the parts of the security infrastructure (firewall, radius server) are kept up-to date by applying all necessary security patches, in order to fix any emerging vulnerabilities of the community infrastructure. Finally, we should periodically confirm application security and eliminate programming faults in the applications running on the web and application servers.

Data protection requires a periodic backup in order to ensure the integrity of the medical data stored in the database. The backup tapes must be kept in a secure off-site location. The disaster recovery plan must be tested periodically by using various disaster scenarios.

3.5 Increase members' trust

The success of a health-related virtual community is based on the frequency and quality of members’ contribution (e.g. medical advices) and on the discreet use of patient sensitive data. Although, patients’ identities can be concealed behind a virtual one, their health record is necessary for the doctor to provide a diagnosis. On the other side, patients should be confident that the identity of the doctor, who receives their data, is valid. This iterative negotiation process [22] assumes that both patient and doctor exchange digital credentials based on the access control policy of each part. Access restriction to sensitive information can be attached to these credentials upon members’ discretion. The community administration authority or any other trusted institution (e.g. the hospital, medical center, ministry of health etc) will be the certificate authority (CA) in this process that guarantees anonymity and atomicity of members at the same time.

Finally, user information from the auditing server can be employed to develop a reputation management application. Patients’ or doctors’ comments on another community member are recorded in the auditing server (Fig.2, 16). A reputation management application, in the application server (Fig.2, 12) will process data and provide each user with a reputation score for any community member based on the community reputation for this member and the direct trust towards this member.

4. COMMUNITY THREATS

This section presents the potential threats for a virtual healthcare community. Threats vary from technical to ethical and regulatory issues. Technical threats that might breach confidentiality and corrupt integrity of medical data or cease availability of healthcare services provided by the community are analyzed in section 4.1. Ethical issues that might arise during operation of the virtual community are mentioned in section 4.2. Section 4.3 provides an overview of the regulatory framework that governs medical data and the rules that must be kept, in order for the medical virtual community to operate in compliance with data protection laws.

4.1 Technical threats

Technical threats target both the information repository and the operational infrastructure of the virtual medical community. A virtual medical community is susceptive to a variety of attacks. Ranging from outside malicious users gaining unauthenticated access, to inside users gaining unauthorized access control to sensitive patient information, all these threats are a major issue that concerns both the CIA (confidentiality, integrity, availability) model and community trust. Identities can be stolen by phising attacks. Denial of service attacks can render the whole community unavailable. Eavesdropping can lead to information leakage, while message disclosure can lead to breach of confidentiality. Web application attacks can damage the database or lead to major information leakage in various ways. The list of threats mentioned previously is indicative and grows as technology advances. Threats are dealt with security measures such as the ones proposed earlier in this paper.

4.2 Ethical

The goal of a virtual healthcare community is mainly to provide patients with medical consultation. The community stores a whole load of medical data in its servers and grants data access to various entities based on their access role and responsibilities. For example, doctors have access to the medical profile of patients in order to provide consultations. If a particular doctor improperly uses patient information to perform genetic or biomedical experiments, or provides medications that violate accepted policies, then important ethical issues arise. Another violation of code ethics would be a patient that harasses other patients. Such behaviors are unethical and raise issues of trust in the community.

4.3 Legal issues

Virtual healthcare communities usually cross national borders and as such, they face several legal issues, such as licensing, accreditation, concerns of identity deception and dependency, which are difficult to be properly addressed by legislative entities.

The opt-out policy adopted by the U.S. Government defines that companies cannot collect consumer’s data if the consumer asks for it. Concerning medical information, U.S. laws [24] assume total confidentiality in several issues (i.e. abortions, contraception or psychological diseases) but delegate decisions to the state laws in others.

European Union has adopted an opt-in model for all personal data, which assumes that all personal information is classified until their owner grants access on them [9]. According to the EC directive for medical data protection (95/46/EC), only health professionals can access medical information and are responsible for protecting confidentiality. According to the Recommendation (97) 5 [8], medical data can be collected without user consent, only for preventing a real danger or in the case of a criminal offence. Moreover, if the law provides for this, data may be collected and processed in order to preserve vital interests of the data subject, or of a third person. In the case of genetic data this includes the members of the data subject's genetic line.

The Greek data protection law [10] is in accordance to the European Directive. Article 2 paragraph b states that medical data are sensitive data. Article 7 paragraph 5d allows the processing of medical data by persons that professionally provide health services but only after permission of the Greek Data Protection Authority (GDPA). Article 6 defines that the data processor needs to ask permission from the GDPA for processing medical data, and the entire data process needs to be analog to the dedicated purpose (article 4). Lastly, article 10 paragraph 3 states clearly that the data processor needs to take appropriate security measures in order to protect privacy of sensitive data.

In the case of a virtual healthcare community, it will abide to the Greek Law, if it operates in Greece and to the law on medical confidentiality (Greek Law, 3418/2005). Thus, the founders and operators of the community need to notify the respective DPA on the data process, specify the purpose of the process, obtain permission on sensitive data and take the appropriate security measures. GDPA requires a security policy, a security plan, a code of conduct and a risk analysis document. International transfer of medical data is governed by article 9 of law no 2472/1997 and is not allowed outside the EU without certain prerequisites.

For each and every international transfer of medical data the competent Data Protection authority must be notified, according to the specific national rules, to, and deem if the transfer is allowed. Possible legal issues that might arise could be the improper use of patient data, selling data to insurance companies and the use of medical data for other than the notified purpose. In case of illegal processing of data, in a virtual healthcare community, if the data processor resides in Greece, then he is subject to penal, civil and administrative penalties according to articles 21-23 of Greek Law 2472/1997. Most EU national laws assume similar penalties.

5. USAGE SCENARIOS

As a first step towards building a concrete evaluation framework, where the proposed architecture will be tested, we carefully design several usage scenarios. These scenarios are expected to expose the architecture's vulnerabilities and can be further utilized in the creation of test cases for a future system that applies the aforementioned architecture. In this section we present several attack scenarios that impact the smooth operation of the community. The security risks usually apply either on the network [15] or on the application level [20]. A holistic security plan also considers attacks "from the inside" of the community [1]. Strictly defined policies and careful auditing will prevent security violations and will provide useful tracking evidence in case of internal attacks or information leaks.

5.1 Scenario 1 - Unauthorized access

5.1.1 Description

Typically the motive behind a hacking attach to the virtual medical community information system is the hacker's interest to prove worthy of gaining access to system and to explore a protected computer system. In a typical scenario, the hacker scans the network of the virtual community and tries to enumerate the community's information infrastructure in order to gather as much data possible concerning the target network. After that, the hacker proceeds with vulnerability scanning aiming to identify open ports and running services that can be exploited and used as points of entry to the system. At the final step, the hacker uses various attacks methods, which aim to exploit identified system vulnerabilities in order to gain access to the community systems. The hacker's ultimate goal is to gain administrative access to an important system server (i.e. the database, the application or the web server), practice on it and subsequently use it to gain access to even more systems.

5.1.2 Critical points

In order to keep the hacker as far from the community servers as possible, it is important to strengthen the fortification of the community infrastructure focusing on the "outside perimeter". The perimeter of the healthcare community comprises the proxy server, where the authentication takes place, the firewall and the wireless access points. These systems must be properly configured in order to prevent unauthenticated user access.

5.1.3 Security mechanisms being activated

The proposed architecture uses authentication tokens, in order to control access. The token mechanism allows encryption of the traffic by combining a used controlled section (i.e. the token that generates the OTP password) and a user id. The potential hacker has to find a way to hijack the encrypted session by exploiting a vulnerability of the authentication mechanism, in order to gain access to the system. Using the proxy as the single point of entry introduces a single point of failure but if set up correctly (with on time patch management) the only danger exists in zero-day exploits.

A second defensive obstacle for the hacker, who manages to hijack the session, will be the firewall, which will identify potential illegitimate traffic and block the attack. The key here is to constantly update the firewall and properly set up the access rules in order to prevent potential breaches. Applying state of the art encryption (e.g. WPA2 with long random passwords or passphrases and the 802.1x protocol) we limit the possibility of a hacker setting a fake access point and gathering enough traffic to hijack a session. Lastly, securing the computer systems of the community members limits the possibility of a 'trojan' or 'worm' attack that will open a backdoor to the community.

5.2 Scenario 2 – Information stealing

5.2.1 Description

Another potential threat is a skillful internal or external user, who sets up a Man in the Middle (MITM) machine and uses it to intercept traffic and steal sensitive medical data. The user sets up a sniffer and listens to ARP packets. In a first step, the attacker learns the IP and MAC addresses of the two communicating parties. Subsequently, the hacker attempts to convince each end of the communication that he is the other end by sending forged ARP packets. As a consequence, all packets pass through the MITM machine and can be processed illegitimately.

A different information stealing attack can be performed in a web-application using for example SQL-injection. In this scenario, the attacker takes advantage of input validation vulnerabilities, queries the database with specially crafted SQL inputs that draw illegitimate information from the database.

5.2.2 Critical points

The critical point in the first scenario is the protection of data exchanged inside the community. Even if the attacker manages to infiltrate the first layer (enter the community perimeter) and perform a MITM attack on the second layer (inside the perimeter), all the intercepted data must be unintelligible and useless to the attacker. The second scenario mainly targets the vulnerabilities of the data exchange applications and mainly refers to the validation of user input.

5.2.3 Security mechanisms being activated

Tokens used for user authentication can also be employed for data encryption. For example, all traffic between an authenticated user and the database server is encrypted using the token and the cryptography algorithm it applies. As a consequence, even if an internal or external attacker manages to become MITM, he will manage to intercept nothing but cryptographic gibberish, which is useless without the decryption key.

Using specialized input filtering modules and a web application firewall in the presentation layer to perform input filtering by methods such as web attack signatures, identification of sql injection characters and dynamic profiling of usual data traffic is the first layer of defense against information stealing attacks such as sql injection. Authorization at the application server level adds an extra obstacle to illegal data access, in case that the malicious user input has not been identified by the web application firewall.

5.3 Scenario 3 – Fake identity

5.3.1 Description

In this case, someone infiltrates the system with a fake identity in order to perform fraud. When the attacker pretends to be a doctor, wrong consultation may be provided to the patients. When pretending to be a patient, then wrong information will be collected.

5.3.2 Critical points

The critical point is to certify the doctor's or patient's identity by using proper authentication methods. Additionally, in case of an identity theft auditing must be in place, in order to prove the fraudulent activity, prevent and undo any damage.

5.3.3 Security mechanisms being activated

The use of tokens as an authentication method along with its registration on a certification authority hinders the possibility of a malicious user entering the community with fraudulent purposes. Additionally, content moderators can function as a second certification authority, preventing the fraudulent users' actions and protecting users from deception. Auditing mechanisms can give an audit trail to the imitator. Finally, reputation mechanism can increase members' awareness on faulty consultation and fraud.

5.4 Scenario 4 – Provide fictional patient data

5.4.1 Description

This attack aims in modifying patient data in the **database**, or infiltrating the mobile sensor network and **transmitting** invalid data. A skillful internal or external user performs the attack by directly sending forged ARP packets to the database, application or web server, intercepting all packets exchanged and altering the medical data in transit. In a similar manner, the attacker can delete information from the database.

5.4.2 Critical points

As in the MITM scenario, the critical point is the protection of the data flow. Even if the attacker succeeds in connecting to the server and performing the attack, the alterations must be recorded and rejected or undone.

5.4.3 Security mechanisms being activated

All mechanisms that confront the MITM attack scenario, such as the authentication and encryption using tokens, will be activated in the first level. In addition to this, backup and transaction auditing mechanisms will allow the detection of data corruption or modification and will assist IT administrators to rollback data in a previous stable state.

5.5 Scenario 5 – System attack

5.5.1 Description

The attack aims to block the smooth operation of the system, in order to obfuscate the community (e.g. in a DOS attack). "*A DOS attack disrupts or completely denies service to legitimate users, networks, systems, or other resources. The intent of such an attack is usually malicious and often takes little skill because the requisite tools are readily available*" [16].   DOS attacks are usually performed through a large number of PCs, infected by Trojans or rootkits, which constitute a botnet network. The attacker (bot-herder) remotely controls the PCs and orchestrates the attack which aims to bring down the entire network. The PCs usually belong to unsuspected users, who are unaware that their computers are infected.

5.5.2 Critical points

The attacks mainly target the application server layer and the designer's aim should be to confront these attacks on the proxy server or the firewall, so that the functionality of the community cannot be impaired. A single vulnerability can be enough to run the system down so a recovery plan should always be ready.

5.5.3 Security mechanisms being activated

The first step in securing the application server is to harden the application firewall by using the latest attack signatures for update. For the attacks that cannot be detected, it is necessary that the firewall and the other network devices are set up correctly to deal with packets arriving at closed ports, with illegitimate packets, etc. Several techniques, such as network ingress filtering, use of BGP to block DOS attacks etc. (see RFC2267 and RFC3382) are tools that administrators can use, in order to fortify their systems. Applying these techniques might cause legitimate traffic to be blocked also, thus the decision to apply them must be taken carefully.

The second step is to set up a patch management suite that operates in a regular basis, in order to limit system's vulnerabilities, which can be exploited in a DOS attack. Lastly, in case of a zero-day exploit, a disaster recovery plan must be able to function properly, in order to bring the soonest possible the community systems back to normal working state.

6. Conclusions

In this paper we presented the desired architecture for a virtual healthcare community. Patients have web access to the community services and provide their medical data using wireless sensor devices and/or web browsers. Doctors access community services either remotely or from inside the hospital. The design of a secure and trustful community is a difficult though interesting task. In this dynamic environment, new applications are added, thus opening new exploits, creating new threats and new attack forms. Security and trust management requires careful handling of all the aforementioned issues and continuous maintenance of the community infrastructure. Our next plans comprise the design of an evaluation plan, which will be based on the presented scenarios and the implementation and evaluation of a prototype application.

7. REFERENCES

[1] Apostolakis, I., Chryssanthou, A., Varlamis, I. 2009. A Holistic Perspective of Security in Health Related Virtual Communities, in Handbook of Research on Distributed Medical Informatics and E-Health, IGI Global.

[2] Becker, M.Y., Fournet, C., Gordon, A.D. 2007. Design and Semantics of a Decentralized Authorization Language. In the 20th IEEE Computer Security Foundations Symposium.

[3] Becker, M.Y., Sewell, P. 2004. Cassandra: Flexible Trust Management Applied to Electronic Health Records. In the 17th IEEE Workshop on Computer Security Foundations.

[4] Blaze, M., Kannan, S., Lee, I., et al. 2009. Dynamic Trust Management. In IEEE Computer Magazine, pp. 42, no. 2, pp. 44 - 52, February 2009.

[5] Demiris, G. 2005. Virtual Communities in Health Care, in StudFuzz 184, 121–122, Springer-Verlag.

[6] Demiris, G., Parker, O.D., Fleming, D., Edison, K. 2004. Hospice Staff Attitudes towards "Telehospice". In American Journal of Hospice and Palliative Care 2004; 21(5): 343-348.

[7] Dorothy, W., Curtis, M.S., Esteban, J. et al. 2008. SMART - An Integrated, Wireless System for Monitoring Unattended Patients. In JAMIA January 2008;15:44-53.

[8] European Council. 1997. Explanatory Memorandum to Recommendation (97) 5 on the Protection of Medical Data.

[9] European Parliament, EU Council. 1995. Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data.

[10] Greek Data Protection Act. 1997. Law 2472/1997, Protection of individuals from sensitive personal data processing.Greece

[11] Jones, V.M., van Halteren, A.T., Dokovski, N.T. et al. 2006. Mobihealth: mobile services for health professionals. In: M-Health Emerging Mobile Health Systems. Istepanian, Laxminarayan & Pattichis (Eds.), Kluwer.

[12] Kui, M., Yue, W., Xu, Z., Xiaochun, X., Gengdu, Z. 2005. A Trust Management Model for Virtual Communities. In the 5th Int. Conf. on Computer and Information Technology.

[13] Laleci, G.B., et al Dogac, A., Olduz, et al. 2008. SAPHIRE: A Multi-Agent System for Remote Healthcare Monitoring through Computerized Clinical Guidelines, in Agent Technology and e-Health, Whitestein. pp. 25-44

[14] Lorincz, K., Malan, D. J., Fulford-Jones, T. R. F. et al. 2004. Sensor networks for emergency response: challenges and opportunities, in IEEE Pervasive Computing;3;16-23.

[15] Maji, A.K., Mukhoty, A., Majumdar, A.K. et al. 2008. Security Analysis and Implementation of web-based telemedicine services with a four-tier-architecture. In proceedings of the 2nd Int. Conf. on Pervasive Computing Technologies for Healthcare, pp 46 – 54.

[16] McClure, S., Scambray, J., Kurtz, G. 2003. Hacking Exposed: Network Security Secrets and Solutions, Fourth edition, 2003, McGraw-Hill/Osborne, pp. 503

[17] Mondy, J., Torresi, M. 2008. CIGNA Creating a Virtual Health Care Community. Article released July 1, 2008 at http://newsroom.cigna.com/

[18] Ng, H.S., Sim, M.L., Tan, C.M. 2006. Security issues of wireless sensor networks in healthcare applications. In BT Technology Journal 24, 2, 138-144.

[19] OASIS. eXtensible Access Control Markup Language (XACML) Version 2.0 core specification, 2005. At www.oasis-open.org/committees/xacml/

[20] Orrin, S. 2004. The twelve most common application level hack attacks, Wathcfire whitepaper. Available at: http://www.emedia.co.uk/FM/GetFile.aspx?id=58740

[21] Schopp, L.H., Hales, J.W. et al. 2004. Design of a Peer-to-Peer Telerehabilitation Model. In Telemedicine Journal and e-Health Jun 2004, 10(2): 243-251

[22] Seamons, K., Winslett, M., Yu, T., Yu, L., Jarvis, R. 2002. Protecting privacy during on-line trust negotiation. 2nd Workshop on Privacy Enhancing Technologies. Springer.

[23] Stanberry, B. 1998. The legal and ethical aspects of telemedicine. Data protection, security and European law. In Journal of Telemedicine and Telecare, 1998;4(1):18-24.

[24] U.S. Department of Health and Human Services, Office for Civil Rights. 2003. Standards for privacy of individually identifiable health information, U.S.

[25] Varlamis, I. and Apostolakis, I. 2007, Self supportive web communities in the service of patients. In proceedings of the Int. Conf. on Web Based Communities, IADIS.

[26] Wang, X., Lao, G., DeMartini, T., Reddy, H., Nguyen, M., Valenzuela, E. 2002. XrML - eXtensible rights Markup Language. In Proc of XML Security. XMLSEC '02. ACM.

[27] Warren, S., Lebak, J., Yao, J. et al. 2005. Interoperability and Security in Wireless Body Area Network Infrastructures. In IEEE-EMBS 2005, pp.3837-38