Chapter 5
A Security Model for Virtual Healthcare Communities

Anargyros Chryssanthou [1], Iraklis Varlamis [2], and Charikleia Latsiou [1]

[1] Auditors Department, Hellenic Data Protection Authority, Ampelokipoi, Athens, Greece, achrysanthou@dpa.gr; clatsiou@dpa.gr

[2] Department of Informatics and Telematics, Harokopio University of Athens, Athens, Greece, varlamis@hua.gr

Abstract Virtual healthcare communities aim in bringing together healthcare professionals and patients in order to further improve the quality of healthcare services and assist healthcare professionals and researchers in their everyday activities. Patient monitoring and medical consultation – the two most popular activities inside virtual healthcare communities – require members' collaboration in a secure and reliable environment. In this environment, patients share their medical data with doctors, expect confidentiality, and demand reliable medical consultation. Apart from a concrete policy framework, 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, giving emphasis on the security issues that arise when attempting to manage risk inside such a community. Following a standardized risk assessment process, which identifies, estimates, and evaluates all potential security risks for the community, a security model is developed, and the community architecture is designed. Finally, a set of usage scenarios, with reference to real events, is employed in order to uncover security risks and illustrate the solutions provided by the proposed architecture.

Keywords Information security • Virtual communities • Health care • ISO 27000 family of standards • Risk assessment

5.1 Introduction

The progress in telecommunication technologies has removed several distance and time barriers and allowed virtual communities to flourish. The members of a virtual healthcare community – doctors, patients, and caregivers – collaborate in order to virtually manage the illnesses and improve the quality of patients' life. Patients submit online requests for advice and share problems and solutions with other patients; doctors cooperate with each other, supervise, and support their patients. Specialized healthcare communities, such as self-supportive patient communities

A.A. Lazakidou (ed.), *Virtual Communities, Social Networks and Collaboration*, Annals of Information Systems 15, DOI 10.1007/978-1-4614-3634-8_5,

that promote peer-to-peer patient communication and medical research communities that support the collaboration of medical professionals, can be assisted by the virtual organizational model.

The most significant issue in healthcare applications is protecting a 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 toward 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 medical data is always available to the rightful 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. Moreover, a set of trust-enabling functionalities, such as transparency, content quality control, and access rights management are desired in order to strengthen users' trust toward the community [14]. As described in [6], 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; and providing a flexible identification mechanism, which preserves anonymity while 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.

A standardized methodology for securing the community infrastructure is preferable to fragmentary solutions that answer specific parts of the problem. The ISO/IEC 27001:2005 standard defines such a methodology for planning and designing an Information Security Management System (ISMS), through a structured process which involves risk assessment, risk treatment planning, selection and implementation of security controls, etc. This structured process is complemented with ISO/IEC 27002:2005, which presents a code of practice for Information Security Management, and, in the e-health sector, with ISO/IEC 27799:2008 [23], which constitutes an ISO/IEC standard directed directly to e-health that provides guide to health information practitioners on how to protect confidentiality, integrity, and availability by implementing ISO/IEC 27002:2005. These security management methodologies combined with ISO/IEC 27005:2008, which is a specialized ISO/IEC standard belonging to the 27000 ISO family of information security standards and deals with information security risk management, are adopted in this current work in order to define a concise security model and architecture.

The contributions of this work are summarized in the following:

- A roadmap for building secure portals for virtual healthcare communities and a holistic approach in developing and maintaining a trustful and secure solution. The first step involves the description of the community structure and the design and development of the initial security model.
- A risk management model, which iteratively identifies and evaluates all the potential hazards and suggests new certification mechanisms, operational policies, and functionalities that can improve the security model, strengthen the protection of the community assets, and increase members' trust.
- A scenario-based validation process, which tests the security model against various incident scenarios that might violate the community processes and lead to loss of confidentiality, availability, or integrity in terms of a data breach on a denial of service attack.
- A study on the legal implications of security incidents involving sensitive medical data in healthcare communities and a presentation of the process followed by the Hellenic Data Protection Authority for auditing, locating, and penalizing violators.

The following section refers to related works that partially cover the need of healthcare communities for security and trust. Section 5.3 gives an overview of the community structure; illustrates the steps of the information security management life cycle, which needs to be followed when developing an Information Security Management System (ISMS); and states the need for a standardized process for developing a holistic security solution. Section 5.4 describes in detail every step of the process that was followed in order to secure the community, from recording community assets and available security controls to identifying threats and potential consequences (technical, ethical, and legal), estimating and evaluating identified risks against selected criteria, and finally forming a risk treatment plan for the community. Section 5.5 presents the community's security model and explains how the available security solutions are orchestrated. Section 5.6 analyzes the specific legal implications behind potential security incidents in a healthcare community and presents the action plan of the Hellenic Data Protection Authority. Section 5.7 presents several potential threats for the community, with reference to real cases, and focuses on the mechanisms of the community's security infrastructure that are activated to confront them. Finally, Sect. 5.8 presents conclusions from this work.

5.2 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 [35] has launched a

virtual community for nutrition and healthcare, which is situated on a Second Life island. The EU-funded project Saphire [28] 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 [12], patient peer support [43], 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 sensor-based monitoring systems have been designed that allow continuous recording of patients' status, such as CodeBlue [32], Scalable Medical Alert Response Technology (SMART) [8], MobiHealth [24], etc. In such dynamic and data-rich environments, a holistic security approach ([2] and [7]) is necessary in order to guarantee confidentiality and reliability and consequently increase community trust. This approach should study the community structure, identify its assets as well as its pros (existing security controls) and cons (existing vulnerabilities), and record potential threats, in order to build a list of potential risks, which can be estimated and evaluated against selected criteria and will potentially lead to possible solutions by means of a well-formed risk treatment plan. A standardized methodology should be employed for this purpose [20–22].

The first step in building a trustful healthcare community is to provide members with a security infrastructure. Patients must be sure that their medical data remain confidential and are constantly available, and their integrity is maintained. Additionally, they must be aware of their virtual caregivers' (doctors, nurses, family or friends etc.) identity and be able to selectively provide access to their sensitive data only to the appropriate right holder. The Health Information Trust Alliance [18] released a security framework for healthcare in March 2009, which is based on well-known standards such as COBIT, NIST, and ISO/IEC 27001:2005 but only available to member organizations subsequent to paying a fee [25]. The framework supports regulations such as the Health Insurance Portability and Accountability Act (HIPAA) [47] and aims in increasing patient confidence in the security of their information. However, it is not clear whether it includes a comprehensive set of privacy principles. In the Cassandra trust management system for medical communities [5], 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 [51] allows the definition of rights and granting policies with validity restrictions. XACML [39] 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) [4] is another XML-based model, which builds on the notion of tunable expressiveness introduced in Cassandra.

Solving the security issues, which relate to the wireless or wired transmission of data [37], and the legal and ethical issues concerning confidentiality of patient data [45] is not always adequate for building trust in the healthcare community. According to [26], trust is subjective, bidirectional 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. In the sections that follow, we present the structure of a healthcare community and its proposed security model and emphasize on the legal aspect of security incidents. This holistic approach and an iterative refinement of the security processes and mechanisms as explained in Sect. 5.4 will guarantee security and increase members' trust to the community.

5.3 A Virtual Community for Monitoring and Tele-Health Care

Virtual communities refer to groups of people that collaborate and discuss their issues, share experiences, consult with experts, and provide and request support by 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 healthcare services at different points in time, in this way bridging geographic distance and time constraints [11, 49].

5.3.1 Community Members

The active members of a virtual healthcare community comprise patients and doctors, as well as people with interest in the community issues, such as patients' family members, researchers, etc. Members have different roles depending on their needs and expertise: patients and family members undertake facilitator roles, while healthcare professionals become moderators for discussion and contents, facilitators and mentors for the community members. The technical 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,

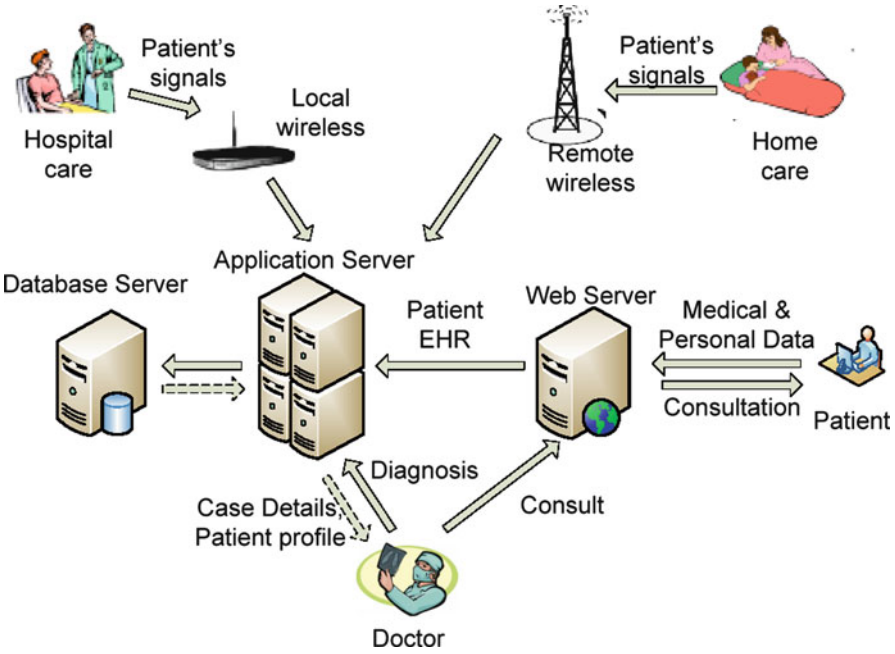

Fig. 5.1 Overview of the community interactions

the employees of the telecommunication service 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.

5.3.2 Community Activities

An overview of the interactions inside the healthcare community is presented in Fig. 5.1. Health status signals are collected using wireless sensors [1, 36] and/or wired devices [46] and are stored in the community servers for future reference and analysis [27]. Patient members are also able to request for advice, diagnosis or treatment suggestion, etc., by using the community portal communication services (e-mail, forum, etc.).

The doctor from inside the hospital is able to access the patient's record (Electronic Health Record – EHR) and make a diagnosis. The doctor replies to patients' requests but also provides consultation based 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.

5.3.3 Securing the Community

Building a security infrastructure is a bottom-up and continuous process. In order to design an effective security infrastructure for any organization, the first step the designer of the organization's Information Security Management System (ISMS) must perform is a detailed and accurate risk assessment of the organization.

A risk assessment according to ISO/IEC 27005:2008 is the overall process of risk analysis and evaluation. Risk analysis comprises of two interconnected stages, risk identification and risk estimation. Risk identification involves identifying assets to the organization, whether human, such as people, technical, such as hardware, information, such as medical data, or critical processes, such as transfers of medical signals. It includes identifying threats, existing security controls, vulnerabilities, and potential consequences. Risk estimation aims to identify a level of risk for each risk identified in the previous process. Risk evaluation comprises evaluating identified risks against selected criteria. These criteria are selected beforehand and could be legal, operational, and organizational.

In order to secure the infrastructure of the community, a specific methodology defined by ISO/IEC 27005:2008 was adopted, which supports the requirements set by ISO/IEC 27001:2005. By using this methodology, assets to the virtual community were located, while emphasis was given to patients and their sensitive medical data. Potential risks were identified by analyzing several reported incidents that relate to the exposure, damage, or loss of patient information. Lastly, these risks were assessed, and a risk treatment plan was formed which leads to selecting the security controls that are incorporated in the proposed community's security model. The community's security model effectively constitutes the community's Information Security Management System, as defined in ISO/IEC 27001:2005.

This methodology is part of a continuous cycle of procedures, which repeats itself as an organization evolves and the legal and operational environments, in which the organization operates, change. Simultaneously, the modus operandi of potential e-crime perpetrators becomes more and more sophisticated as technology evolves and newer tools are developed, which make it easier for a perpetrator to steal valuable data. This cycle consists of four stages, namely, Plan, Do, Check, and Act, as depicted in Fig. 5.2.

An Information Security Management System is a constantly evolving part of an organization, which has to be carefully designed and implemented in the Plan and Do phases. However effective the ISMS of the organization is, the implementer will have failed if he does not complete the cycle over and over again by checking the ISMS' effectiveness and correcting any identified security issues. Integral part of this cycle is the risk management methodology that is being followed by the implementer with its key product being the risk treatment plan. The latter, in the case of a medical community, produces a security model, which is the core of the community's ISMS. This security model is the key contribution of this chapter and is depicted in Sect. 5.5.

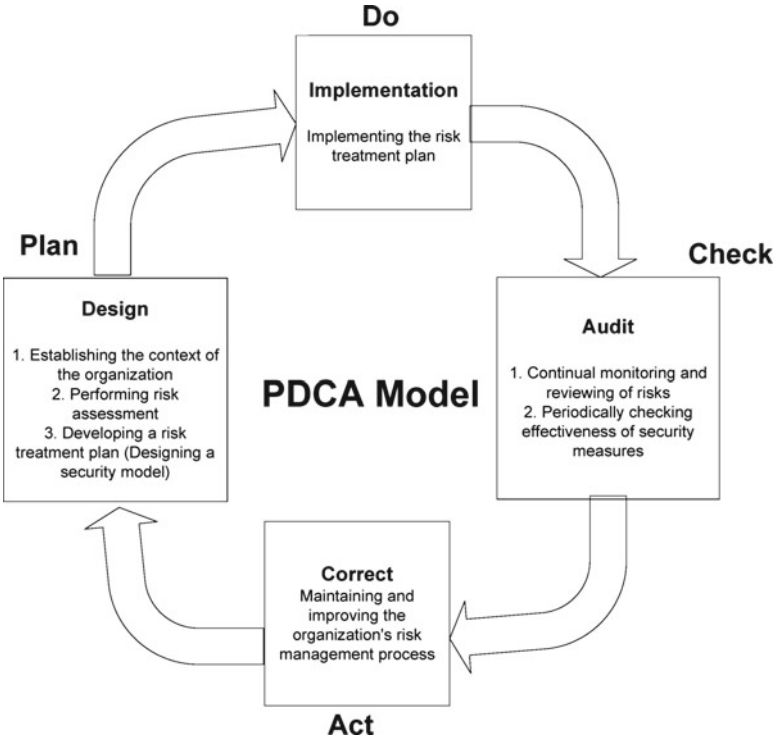

Fig. 5.2 The information security management cycle (combining ISO/IEC 27001:2005 and ISO/IEC 27005:2008)

5.4 The Risk Management Model

5.4.1 Risk Identification

The first step in the risk assessment model applied to the community, as defined by ISO/IEC 27005:2008, is risk identification which comprises the recording of (a) valuable community assets, (b) potential community threats, (c) existing security controls, (d) detected vulnerabilities, and (e) consequences of potential incident scenarios in regard to CIA (confidentiality, integrity, availability).

5.4.1.1 Community Assets

Identified assets are the active members of the community, medical data, and the internal structure of the community, which consists of information systems and physical premises. Assets need to be protected in the context of the traditional CIA

model of security with CIA standing for confidentiality, integrity, and availability. Data must remain confidential, maintain their integrity, and remain constantly available. The same applies to systems. A distributed denial of service attack, for example, could render the whole community systems unavailable to patients and doctors, which could even lead to loss of human life in case of an emergency due to unavailability of a patient's electronic health record. Thus, assets need to be assessed in the context of potential threats, which are mainly technical, and consequences, which can be ethical and legal.

5.4.1.2 Threats for the Community

The threats for a virtual healthcare community are mainly technical, and their basic source is the human factor: malicious attackers, naïve users and administrators, malicious insiders, and security-unaware users. These threats have consequences that relate to potential incident scenarios which are represented in Sect. 5.4.1.5. These consequences are multifold with aspects covering technical to ethical and regulatory issues.

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 phishing 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. Data breaches, security incidents that include theft, or loss of digital media, such as USB sticks containing an organization's data, constitute a phenomenon that multiplies nowadays and can have serious consequences for a virtual community.

The list of threats mentioned previously is indicative and grows as technology advances. Threats are dealt with security measures such as the ones proposed later in this chapter.

5.4.1.3 Security Controls

The identification of assets and threats is followed by a careful recording of the community shields. The security controls, which are readily available, should also be recorded. This will inform us on the state of the existing controls. The list of potential security controls contains physical, technical, and administrative solutions such as security guards, surveillance cameras and locked doors, data encryption, smart cards, network authentication, access control lists (ACLs), file integrity auditing software, security training programs, and disaster recovery plans.

In order to identify existing controls, one should read the previous information security management records, interview the people who are responsible for information security, and perform an internal on-site audit of existing security infrastructure. All the available controls will be evaluated in the following steps, depending on their cost, vulnerabilities, and application status, as well as on the criticality of the processes they guard.

5.4.1.4 Identifying Vulnerabilities

A threat can become a real danger for the community if it manages to exploit existing vulnerabilities. For this reason, the identification of vulnerabilities is an important step of the risk management procedure and should be performed in a constant basis.

The vulnerabilities may be found in the human factor (personnel), in organizational and management routines, in the technological infrastructure, and in the physical premises of the organization. The absence of an access control policy, the sharing of passwords among different users, and the lack of a secure data transmission method are some examples of vulnerabilities. Although, the existence of vulnerabilities does not necessarily mean harm, all identified weaknesses should be properly treated or else should be constantly monitored.

The answer to vulnerabilities is implementing security controls that minimize the vulnerability's effect or totally eradicate the vulnerability. For example, a patch for a vulnerability, which resides in the operating system, will fix the vulnerability and immunize the system to a potential exploit against the vulnerability in question. Additionally, one should also check whether the existing controls effectively cover the weaknesses of the system or whether they should be combined or enhanced to improve a system's defense.

5.4.1.5 Identifying Consequences

Technical threats that might breach confidentiality and corrupt integrity of medical data or cease availability of healthcare services provided by the community were analyzed in Sect. 5.4.1.3. Examining these threats, as part of our risk identification procedure, means constructing incident scenarios and identifying their consequences. These consequences will certainly include technical issues, as new security measures will have to be implemented as countermeasures to these incident scenarios; operational cost, as these security measures are not free; ethical issues depending on the nature of each incident scenario; and legal issues depending on the legal and regulatory boundaries that were taken into account when defining the ISMS policy for the virtual community (clause 4.2.1b in ISO/IEC 27001:2005). In the next two sections, mainly incident scenarios, which lead to ethical issues, will be presented (Sect. 5.4.1.5.1) as well as 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 (Sect. 5.4.1.5.2).

Ethical Issues

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. The code of ethics would also be violated, in a different case scenario, by a patient who harasses other patients. Such behaviors are unethical and raise issues of trust in the community.

Legal Issues

Events that raise ethical issues usually include legal penalties. A data breach incident, including, for example, use of data for other than notified purpose, is subject to civil, administrative, and penal sanctions, which are imposed, among others, from data protection laws. 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.

In the hypothetical scenario, where the administrators[1] of the virtual community illegally process medical data, they could receive administrative penalties (e.g., monetary fine) and will be obliged to modify the way they process medical data. In case of a data breach, once again, the community's administrators will be subjected to an administrative control and later, based on the control's findings, receive a monetary fine and be responsible to adjust the community's security measures, in order to avoid similar data breaches in the future.

Since the implications from a data breach incident or a denial of healthcare service may be fatal, it is important that the community's administrators are familiar with the legal consequences that might arise for them in such unfortunate incident scenarios. For this reason, the legal framework that protects members of medical and healthcare communities in respect of handling sensitive medical data and potential security violations is explained in a separate section (Sect. 5.6).

[1] The administrators of the community could be physical entities (persons) or legal entities, in terms of an authority which supervises the community. A hospital could be the authority in question.

5.4.2 Risk Estimation

Having identified threats and potential consequences, the next step consists of estimating the identified risks. Risk estimation aims to identify level of risk for each risk identified in the previous process (risk identification).

A distributed denial of service attack would be rated as severe in regard to business impact, due to loss of availability, but low in regard to likelihood of event as attackers would not normally aim to bring a medical community down. An accidental loss of community data contained in an USB stick on the other hand would be rated as severe in regard to business impact, due to loss of confidentiality, and medium to severe in regard to likelihood of event. Statistics show that data losses through lost media included 42,212,702 sensitive records with a mean of 659,573 records per incident during the last 5 years, while 3% of data loss incidents in 2008 were due to lost media. However, only 12% of these incidents concerned the medical sector [10].

5.4.3 Risk Evaluation

Risk evaluation comprises evaluating identified risks against selected criteria. These criteria are selected beforehand and could be legal, operational, and organizational. For the case of the proposed virtual community, selected criteria range from loss of human life and legal consequences to likelihood of event, with the latter one balancing the former ones.

Evaluating the above-mentioned two examples would lead to a medium level of risk for the distributed denial of service attack scenario due to low likelihood of event and a high level of risk for the data loss scenario due to the likelihood of event and the potential consequences, which can include even loss of life.

Completing the risk evaluation, existing identified risks have been successfully assessed. The next stage is forming a risk treatment plan based on the conducted risk assessment.

5.4.4 Risk Treatment Plan

Upon completion of a risk assessment, a risk treatment plan must be formed (clause 9 in ISO/IEC 27005: 2008). The plan must explain the selection of security controls in relation to identified risks.

An indicative list of security controls, which can be used to reinforce the community's security infrastructure, is listed in the following table (Table 5.1).[2]

[2] A complete list of security controls that are applicable to any kind of organization structure can be found in ISO 27001:2005 (Annex A) and in ISO 27002:2005.

Table 5.1 Selected security controls from ISO/IEC 27001:2005

Selected control	ISO/IEC 27001:2005 clause	Implemented measure
Access control	A.11	Overall access and behavior policy, all the applicable security measures in this family of controls
Audit logging	A.10.10.1	Audit logs of application and systems
Management of removable media	A.10.7.1	Endpoint security solution
Business continuity management	A.14.1	Disaster recovery plan, periodic security checks (physical and logical security, application and hardware layer)
Information security policy document	A.5.1.1 (and clause 7.2 in ISO/IEC 27799:2008)	Overall access and behavior policy
Review of the information security policy document	A.5.1.2	Policy reviewed at planned intervals of after periodic security checks
Allocation of information security responsibilities	A.6.1.3	Overall access and behavior policy, definition of community roles
Controls against malicious code	A.10.4.1	Filtering module, web application firewalls, signatures of known attacks for routers and firewalls, antivirus software (constantly updatable)
Network controls	A.10.6.1	Firewall, web application firewall, 802.1x wireless access, EAP-TLS authentication, digital certificates
Monitoring system use	A.10.10.2 (and clause 7.7.10.2 in ISO/IEC 27799:2008)	Audit logs of application and systems
Protection of log information	A.10.10.3	Encrypted log storage, two-entity access
Administrator and operator logs	A.10.10.4	Logging of administrator's access
Input data validation	A.12.2.1	Filtering module, web application firewalls
Cryptographic controls	A.12.3	Encrypted hard drives (read in the presence of hardware tokens), authentication through e-token devices
Information backup	A.10.5.1 (and clause 7.7.7.2 in ISO/IEC 27799:2008)	Periodic backup, keeping of backup tapes in separate locations
Technical vulnerabilities management	A.12.6.1	Regular updates to all devices and applications by means of a patch management suite
Data protection and privacy of personal information	A.15.1.4	Compliance to applicable data protection laws
Physical and environmental security	A.9	Supply generators, physical access only to authorized personnel (with combined use of access cards and access code)

Based on the selected controls, a security model for the community is formulated. Along with this security model, the architecture of the community's Information Security Management System is designed. This architecture is explained in the following section (Sect. 5.5).

5.5 Multitiered Security Model

5.5.1 Architecture

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, it must be certified 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 and 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 [33] 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. 5.3, point 3). Moreover, typical web attacks such as cross-site scripting or SQL injection, which may be performed even 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 multilayer 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.

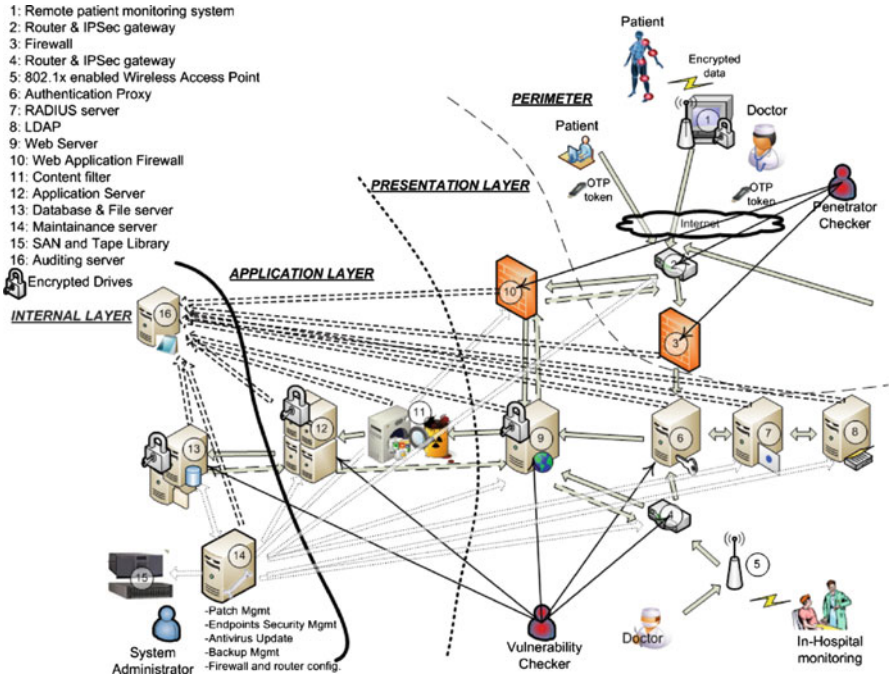

Fig. 5.3 Security-enabled architecture

Figure 5.3 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.

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. 5.3, point 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 patient monitoring systems (Fig. 5.3, point 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 [52]. The encrypted signals, which are collected to the base of the remote monitoring system, can be forwarded via the secure router (Fig. 5.3, point 2) to the application server. Devices that do not support data encryption must be cable-connected to the base.

5.5.1.1 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 in the application layer (Fig. 5.3, point 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. 5.3, point 9) do not attack the application server (Fig. 5.3, point 12). Commercial web application firewalls (Fig. 5.3, point 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. 5.3, dotted arrows) using signatures of known attacks, provided by the firewall's manufacturer, and create additional signatures for illegitimate traffic.

5.5.1.2 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. 5.3, point 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, user roles can be distributed by the authorities of the community, and access credentials can be strengthened by asking users to log in with e-token devices and passwords. The distribution process will be performed by using appropriate registration forms, which will be processed by the community authorities, in order for the users' accounts to be opened and the corresponding credentials to be assigned. The configuration of the application server is performed in the maintenance server (Fig. 5.3, point 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 [4] 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. In special cases, where the normal security protocols need to be overridden due to healthcare priorities, then an access policy exception is necessary, as stated in ISO/IEC 27799:2008 (clause 7.2 in ISO/IEC 27799:2008). Such exceptions will be later incorporated in the access policy provided that they constitute a repeated scenario in everyday practice, which was not predicted during the initial design of the access policy, and not a one-time incident. The provision for such exceptions will be incorporated in the community's access policy.

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 log-in procedures, password quality, and privacy rights to user roles and credentials. It will also include provisions for cases of access policy exception scenarios, which, as mentioned earlier, cause an override of the normal access control mechanisms and roles. 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' consent, an extra permission from the Hellenic Data Protection Authority (Greek Data Protection Law 2472/1997, Article 7, par. 2 [29]). An access policy, a security plan, a Code of Deontology, and a risk analysis document are required by the Hellenic Data Protection Authority in order to grant permission for processing sensitive data.

5.5.1.3 Community Internal Network

In order to improve community trust, firstly, members' responsibilities must be defined, and consequently, it must be certified that members carry them out properly. Auditing can assist in this direction. Every single user action, either local or remote, must be logged (Fig. 5.3, dashed arrows). Logs must be securely stored, in an encrypted format, for a period of time, depending on the criticality of the application that the user uses. This criticality derives from the type of data the application processes and the type of function it performs. Access to the logs will be allowed only in presence of at least two administrative entities of the community with different roles (e.g., 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 as well as retaining the former content of data (in case of insert, update, delete actions) and the associated audit record (as mandated in clause 7.7.10.2 of ISO/IEC 27799:2008)

will allow back tracing to the perpetrator in case of any improper data access and will also allow reconstructing the previous state of the patient's medical record.

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. A secure destruction policy, including provisions for secure disposal of media (used hard drives, tapes) will be in place to ensure that health information is securely destroyed when no longer required for use.

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

5.5.2 Security Maintenance Processes

As a secure infrastructure is important for the operation of the virtual healthcare community, it must be ensured 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, application security should be periodically confirmed and programming faults in the applications running on the web and application servers must be eliminated.

A periodic check of security might lead to changes to the overall access and behavioral policy of the community. This policy must also be reviewed at planned intervals, in order for it to remain effective.

Data protection requires a periodic backup in order to ensure the integrity of the medical data stored in the database. The community's data are backed up in an encrypted format and stored in backup tapes, which are kept in a physical secure off-site location (clause 7.7.7.2 in ISO/IEC 27799:2008). Finally, the disaster recovery plan must be tested periodically by using various disaster scenarios.

5.5.3 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 [44] 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 support a reputation management application. Patients' or doctors' comments on another community member are recorded in the auditing server (Fig. 5.3, point 16). A reputation management application, in the application server (Fig. 5.3, point 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 toward this member.

5.5.4 Physical Security

Having adopted all these previously analyzed security mechanisms in the community's security infrastructure model, the community has achieved a sufficient level of logical security. By means of the security measures analyzed earlier, the community is protected from any attack that does not require physical presence.

However, the case should be considered where an attacker gains physical access to a server and launches an attack, such as a cold-boot attack against a server with encrypted data. In order to avoid such a scenario, physical security needs to be in place as well. Valuable IT equipment, according to ISO/IEC 27002:2005, "should be physically protected against malicious or accidental damage or loss, overheating, loss of mains power etc." Thus, the community's information systems should be protected by supply generators. Moreover, physical access to the information systems' room should be allowed only to authorize IT staff (administrators) or personnel accompanied by authorized IT staff and only with the combined use of access cards and access code. Furthermore, the hardware tokens needed to access the community servers' encrypted hard drives are also kept in a secure location, elsewhere than the community's IT room. Monitoring devices, security guards, and other security measures are advisable and subject to the importance of the application.

5.6 Legal Parameters

The following paragraphs present how data protection laws are applied to personal data or to cases where appropriate security measures are not implemented. Emphasis is given to medical data, since these are the data being processed inside a healthcare

community. First, the legal and regulatory framework is presented. The framework applies to medical data, whether this data resides in EU, in the United States, or is subject to international transfers. Then the potential legal consequences are illustrated by means of a paradigm, which is based on the assumption that the proposed virtual community operates in Greece. This paradigm presents how the Greek Data Protection Law (Law 2472/1997) can be employed in favor of the community.

5.6.1 Legal and Regulatory Framework

The opt-out policy adopted by the US Government defines that companies cannot collect consumer's data if the consumer asks for it. Concerning medical information, US laws [48] 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 [15]. According to the EC directive on the protection of individuals with regard to the processing of personal data and on the free movement of such data (95/46/EC), only health professionals can access medical information and are responsible for protecting confidentiality. According to the Recommendation (97) 5 [16], medical data can be collected without user consent, only for preventing a real danger or in the case of a criminal offense. 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 (Law 2472/1997) is in accordance to the European Data Protection Directive (95/46/EC). Article 2 paragraph b states that medical data are sensitive data. Article 4 states that in order for personal data to be lawfully processed the entire data process needs to be analog to the dedicated purpose. Article 6 defines that the "data controller"[3] needs to notify the Hellenic Data Protection Authority (HDPA) on processing personal data and describes in short what this notification should include. Article 7 paragraph 1d allows the processing of medical data, subsequent to the HDPA's permission, by persons that professionally provide health services and are subject to the duty of confidentiality or to a

[3] The term data controller is analyzed in detail in Opinion 1/2010 of Article 29 Working Party [3] (Working Party that was set up under Article 29 of EU Data Protection Directive) on the concepts of "controller" and "processor." This opinion is available on http://ec.europa.eu/justice_home/fsj/privacy/workinggroup/wpdocs/2010_en.htm (Last accessed on June 4, 2010). According to Article 2 paragraph d' of the EU Data Protection Directive, the term "data controller" "shall mean the natural or legal person, public authority, agency or any other body which alone or jointly with others determines the purposes and means of the processing of personal data." The same definition for the term "data controller" is given in ISO 22857:2004 (Terms and Definitions, paragraph 3.3).

Code of Deontology. This processing is allowed under the condition that it is necessary for medical prevention, diagnosis, care, or management of health care services. With the amendment of the law in 2006 (Law 3471/2006) [31], Article 7A paragraph 2d adds some exemptions to the previous-mentioned article deriving some data controllers from the notification and permission process. However, these exemptions do not include legal entities and organizations providing medical services, as well as data controllers that collect and processes medical data as part of a telemedicine project or in order to provide medical services through a computer network. Lastly, Article 10 paragraph 3 states clearly that the data controller needs to take appropriate security measures in order to protect privacy of sensitive data. The HDPA requires a security policy, a security plan, a Code of Deontology, a risk analysis document, and a secure destruction policy. In the case of a virtual healthcare community, if it operates in Greece, it will abide to general sanctions of Greek Law and to the Greek Data Protection Law (Law 2472/1997) and the Greek Medical Code of Deontology (Law 3418/2005 [30]) in particular.

In general, data controllers can follow during the transfer of personal health information (PHI) the ISO/CD 22857:2004 standard, which provides guidelines on data protection to facilitate transborder flows of PHI. This specialized ISO standard does not in any case replace national legislations on data protection, while "*in cases, where a multilateral treaty between a number of countries has been agreed (e.g. the EU Data Protection Directive), the terms of that treaty will take precedence"* [19]. As far as Greece is concerned, international transfer of medical data is governed by Article 9 of the Greek Data Protection Law (Law 2472/1997) and is not allowed outside the EU without certain prerequisites. For every cross-border exchange of medical data, the competent Data Protection Authority must be notified. The authority must examine the specific national rules and determine whether the data transfer is allowed or not. Possible legal issues that might arise could involve improper use of patient data, selling data to insurance companies, and use of medical data for other than the notified purpose. In case of illegal processing of data, if the data controller resides in Greece, then he is subject to penal, civil, and administrative sanctions according to Articles 21–23 of Greek Data Protection Law (Law 2472/1997). Most EU national laws assume similar penalties.

5.6.2 Application of Greek Data Protection Law to a Virtual Healthcare Community

Supposing that the proposed virtual community resides in Greece, it will abide to several laws, with Greek Data Protection Law being the most significant. The community, being a data controller for sensitive medical data, does not belong to the exemptions of Article 7A paragraph 1d, as explained earlier, because it collects and processes medical data in order to provide telemedicine services. Thus, the data administrator of the community is obliged to notify the HDPA on processing sensitive medical data and request permission for this process. The notification process is depicted in Fig. 5.4.

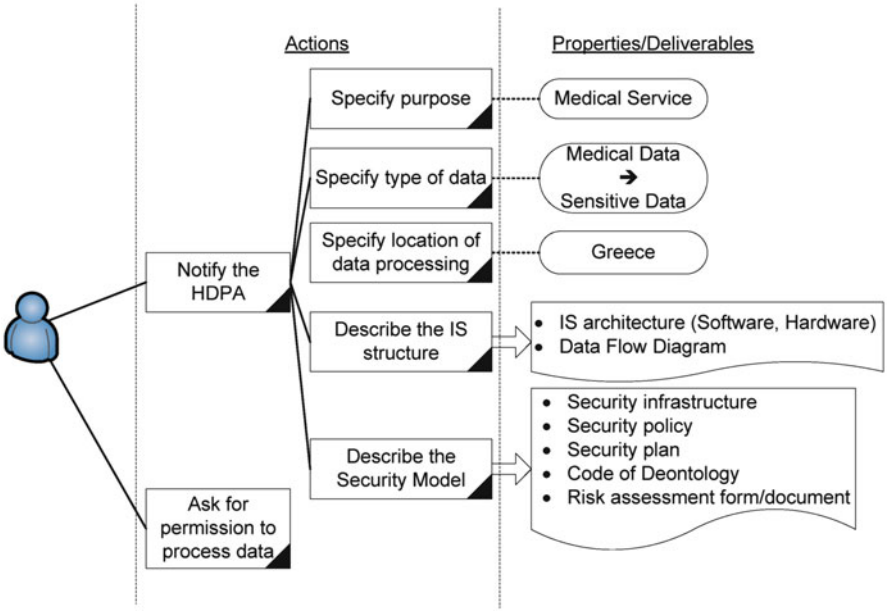

Fig. 5.4 The HDPA notification process

The administrator(s) of the community must notify the HDPA that they process medical data, which are sensitive data, for the purpose of providing medical services, and that they reside in Greece. The administrators must also inform the HDPA on the architecture of their information system and the data flow in it. Finally, they should submit a detailed security report, which contains the security infrastructure of the community, the security policy, the security plan, the Code of Deontology, and the latest performed risk assessment. They can also optionally submit their secure destruction policy document, which is not obligatory in the HDPA notification process. With all these completed, the data controller (i.e., the virtual community operating authority) has completed its notification submission and waits for it to be examined and approved.

In Fig. 5.5, the HDPA notification examination process is depicted. The HDPA auditors examine the type of data processing and the type of processed data; they check the community purpose and examine the information systems and the security model of the community portal. If something is missing, they ask for clarifications or even modifications of the original process in an iterative process that improves information security and enhances the achieved level of data protection. If everything is ok, the permission to process data is granted to the community administrators. If the administrators deny or do not manage to comply with the HDPA requirements, then the permission to process data is denied.

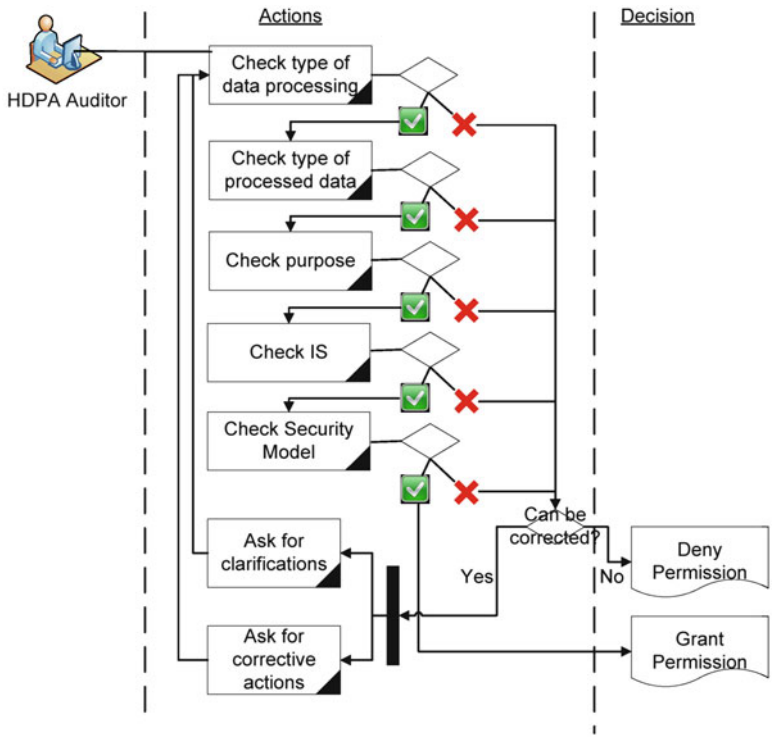

Fig. 5.5 The HDPA notification examination process

5.7 Usage Scenarios

As explained in Sect. 5.3, the development of a secure solution for virtual healthcare communities comprises design, implementation, assessment, and redesign steps is an iterative process. In the following paragraphs, we present several usage scenarios, which aim to expose the vulnerabilities of the model and provide future community developers with a set of test cases for the evaluation of their implementations. Each scenario contains a case description, a summary of the community assets which are on risk, a list of the security mechanisms that will be activated, and a reference to real-world incidents that justifies the importance of each case. The security risks usually apply either on the network [33] or on the application level [38]. A holistic security plan also considers attacks "from the inside" of the community [2]. Strictly defined policies and careful auditing will prevent security violations and will provide useful "tracking" evidence in case of internal attacks or information leaks.

5.7.1 Scenario 1: Unauthorized Access

5.7.1.1 Description

Typically the motive behind a hacking attack to the virtual medical community information systems 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.

5.7.1.2 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.

5.7.1.3 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 user 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), the possibility of a hacker setting a fake access point and gathering enough

traffic to hijack a session is reduced. Lastly, securing the computer systems of the community members limits the possibility of a "Trojan Horse" or "worm" attack that will open a backdoor to the community.

5.7.1.4 Real-World Incident

On February 2006, the University of Washington Medical Center reports an intrusion incident,[4] during which hackers gained access for almost 18 months to more than 200 computers, housing medical and business records. The breach occurred because someone failed to install security patches. The solution was to remove the breached computers from the network and wipe clean their hard drives. As a final step, a commercial intrusion-alert system was installed in all machines.

5.7.2 Scenario 2: Information Stealing

5.7.2.1 Description

Another potential threat is a skillful internal or external user, who steals sensitive information in order to perform fraud, identity theft, etc., and achieve personal profit. In the simplest case, an individual internal user, who has physical access to the system, can copy data in a portable drive (user A in Fig. 5.6).

A different information stealing attack can be performed in a web application using 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 (user B in Fig. 5.6).

Another frequent incident is hardware theft or accidental loss of hardware (i.e., laptops, portable storage devices, or backup media) which results in loss of medical information (user C in Fig. 5.6).

5.7.2.2 Critical Points

The critical point in the first scenario is the protection of data stored inside the community's servers. Since the attacker has already managed to infiltrate the first layer (enter the community perimeter) by means of physical presence, he/she should not be able to retrieve and store any sensitive medical data in a portable drive. The second scenario mainly targets the vulnerabilities of the data exchange applications and mainly refers to the validation of user input. The third scenario

[4] DataLoss database report of incident 224 [cited June 4, 2010]. Available from http://datalossdb.org/incidents/224-hacker-gains-opportunity-to-view-patient-medical-records

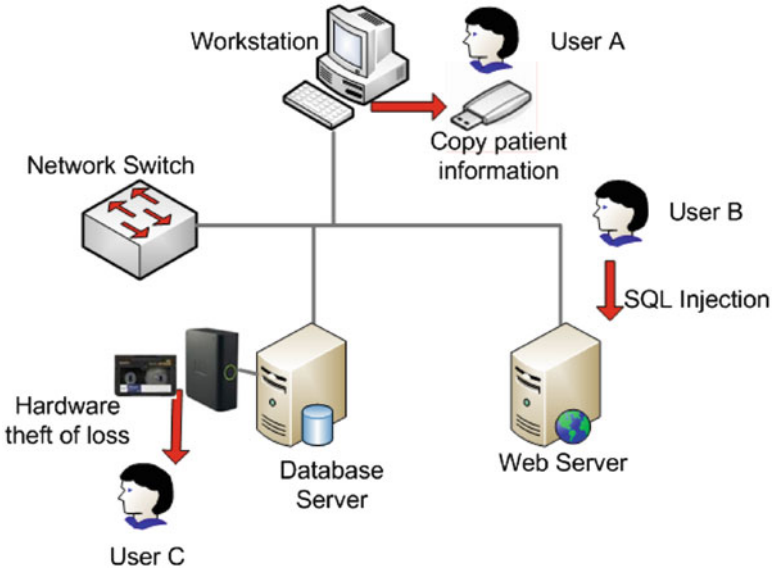

Fig. 5.6 Information stealing

exploits vulnerabilities in the physical layer (e.g., the potential absence of an access control system protecting the IT room) as well as the logical layer (e.g., the existence of unencrypted hard drives containing sensitive data).

5.7.2.3 Security Mechanisms Being Activated

In the first scenario, having an endpoint security solution in place, the community is able to block or control access to portable media devices by its users. Thus, the attacker either cannot use a portable media device in the first place or cannot copy any sensitive data to a portable drive, since the endpoint security solution policy prohibits it. Either way such an attacker is not able to steal any data from the community.

In the second scenario, 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.

In the third scenario, encrypting the server's hard drives and keeping the hardware tokens needed to read them in a separate location harden the attacker from stealing a server hard drive containing sensible data. The same applies to backup

tapes, which are encrypted and thus effectively useless to an attacker. Additionally, the existence of physical security measures (access card systems, access codes, security cards, etc.) makes it even harder to access the community's IT room or steal a backup tape from its physically secure storage location.

5.7.2.4 Real-World Incident

On February 2009, St. Anthony Central Hospital in Denver reported that for 18 months, an employee was stealing records by means of a USB drive (as many as 20 records per week) and used them to make fake driver's licenses and counterfeit Social Security cards.[5]

The Royal Bolton Hospital reported on January 2008 the theft of a computer that contained the private details of 350 chest patients.[6] The hospital contacted all patients to inform them of the theft, but insisted that all information is data-protected and cannot be accessed by anyone other than the relevant hospital staff. In order to improve security, the hospital recalled all its computers and laptops so that vital security software can be installed, which will encrypt patients' details. Additionally, encryption software was installed on all memory sticks and pen drives. Finally, all information was planned to be transferred to a central server and hosted in a secure storage network – rather than on individual hard drives.

5.7.3 Scenario 3: Fake Identity

5.7.3.1 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.

5.7.3.2 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 and prevent and undo any damage.

---

[5] DataLoss database report of incident 2143 [cited June 4, 2010]. Available from http://datalossdb.org/incidents/2143-employee-steals-patients-medical-records-for-counterfeiter

[6] DataLoss database report of incident 1935 [cited June 4, 2010]. Available from http://datalossdb.org/incidents/1935-laptop-containing-the-personal-details-of-about-200-cancer-patients-stolen

5.7.3.3 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 mechanisms can increase members' awareness on faulty consultation and fraud.

5.7.3.4 Real-World Incident

The DataLoss Database and the World Privacy Forum report several incidents on identity theft (e.g., [13]), which resulted in fraudulent social security numbers and abuse of insurance company benefits. The motive behind a fake or stolen medical identity is to obtain medical services, goods, or money by falsifying claims for medical services and falsifying medical records to support those claims.

However, the danger can be even greater, when the theft is discovered during the course of a medical emergency. In this case, the medical profile is distorted by several false entries added by the criminal. This was the case of a Florida woman who discovered that someone impersonating her had caused false entries to be placed in her medical file as reported by the Federal Trade Commission [17]. In the case of a virtual community, the threat of fake identities is bigger, since patients may not have personal contact with their doctors.

5.7.4 Scenario 4: Provide Fictional Patient Data

5.7.4.1 Description

This attack aims in modifying patient data in the *database*, or infiltrating the mobile sensor network and *transmitting* invalid data.

The attacker (MITM – Attacker in Fig. 5.7) gains physical access to the community's premises and sets up a Man in the Middle (MITM) machine and uses it to intercept traffic and steal sensitive medical data. In a first step, the attacker sets up a sniffer and listens to ARP packets. The attacker's initial objective is to learn the IP and MAC addresses of the two communicating parties (Workstation X, Workstation Y, in Fig. 5.7). Workstation X sends an ARP request message to Workstation Y asking the MAC address of Workstation Y. Workstation Y replies with an ARP reply message. The attacker sniffs both messages and subsequently sends a forged ARP reply message (Arp Attack 1, Fig. 5.7) to the requestor's machine (Workstation X) saying "I own the IP you have requested. My MAC address is '00-0G-7E-3T-4C-98.'" Workstation X will gladly update its ARP table with the new one that it has just

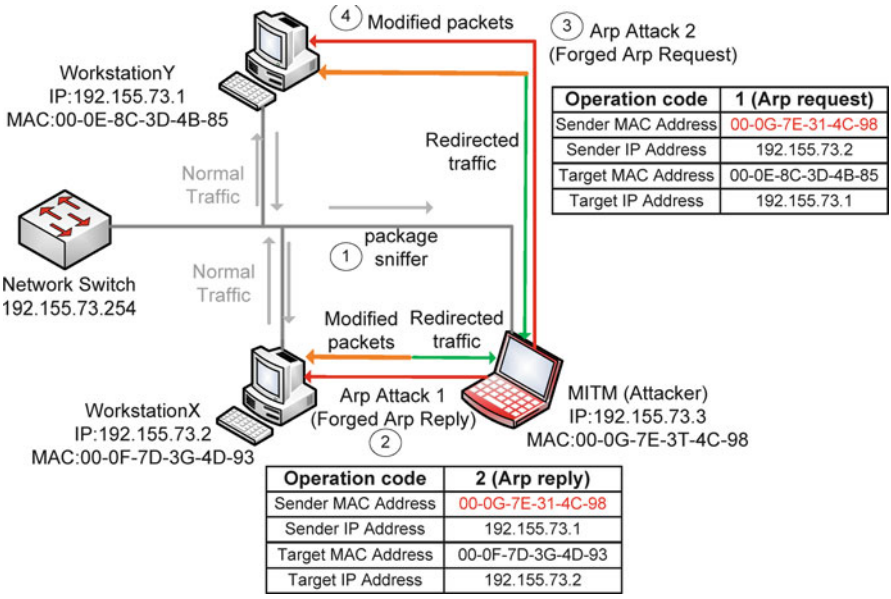

Fig. 5.7 A tampering man in the middle attack

received. At the same time, the attacker sends a forged ARP request message to Workstation Y (Arp Attack 2, Fig. 5.7) saying that he owns the IP address of Workstation X. As a consequence, all packets pass through the MITM machine and can be processed illegitimately. The attacker is able to intercept all packets exchanged, alter the medical data in transit, and delete information from the database.

5.7.4.2 Critical Points

In the MITM scenario, the critical point is the protection of the community network's infrastructure. Even if the attacker succeeds in gaining physical access to a community's network port, he must not be able to gain actual network access. Even if the attacker manages to obtain an actual network IP address, he must not be able to intercept or alter legitimate traffic.

5.7.4.3 Security Mechanisms Being Activated

Several mechanisms are activated in the first level in order to confront the MITM attack. The first step is having employed physical security measures to deter the malicious attacker from gaining access to an actual network port. Even if the attacker

obtains the latter, his presence in the network will be audited by the community's auditing mechanisms. Since the community's network infrastructure is carefully configured, having possibly static routes configured in respect to critical systems, the attacker will not be able to fool any user that the MITM is a critical application server. If the attacker manages to infiltrate by means of a MITM attack into the communication between two legitimate users, the intercepted traffic will be gibberish, since the traffic is encrypted by means of tokens. In case that all before-mentioned measures fail, 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.

5.7.4.4 Real-World Incident

Although MITM attacks mainly target banks, one should expect such attacks in a medical community. The aim of these attacks is to steal passwords, pin, social security numbers, etc. A typical example is the exploit found on SiteKey, a security control used by several web banking applications, which allowed criminals to imitate the original web site and steal data from clients.[7] In a more recent example, Kevin Mitnick set up a Man in The Middle Server that intercepted his call to the IVR of Washington Mutual Bank and managed to steal his account number and four digits of his social security number.[8]

5.7.5 Scenario 5: System Attack

5.7.5.1 Description

The attack aims to block the smooth operation of the system, in order to obfuscate the community (e.g., in a DDOS attack). "A DDOS 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" [34]. DDOS attacks are usually performed through a large number of PCs, infected by Trojan Horses 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.

[7] First live SiteKey exploit seen in operation (30/10/2007) [cited June 4, 2010]. Available from http://cr-labs.com/publications/SiteKeyExploit-20071030–1.pdf

[8] Former Hacker Tackles IVR and Voice Biometric Security [cited June 4, 2010]. Available from http://www.speechtechmag.com/Articles/Editorial/FYI/Former-Hacker-Tackles-IVR-and-Voice-Biometric-Security-50358.aspx

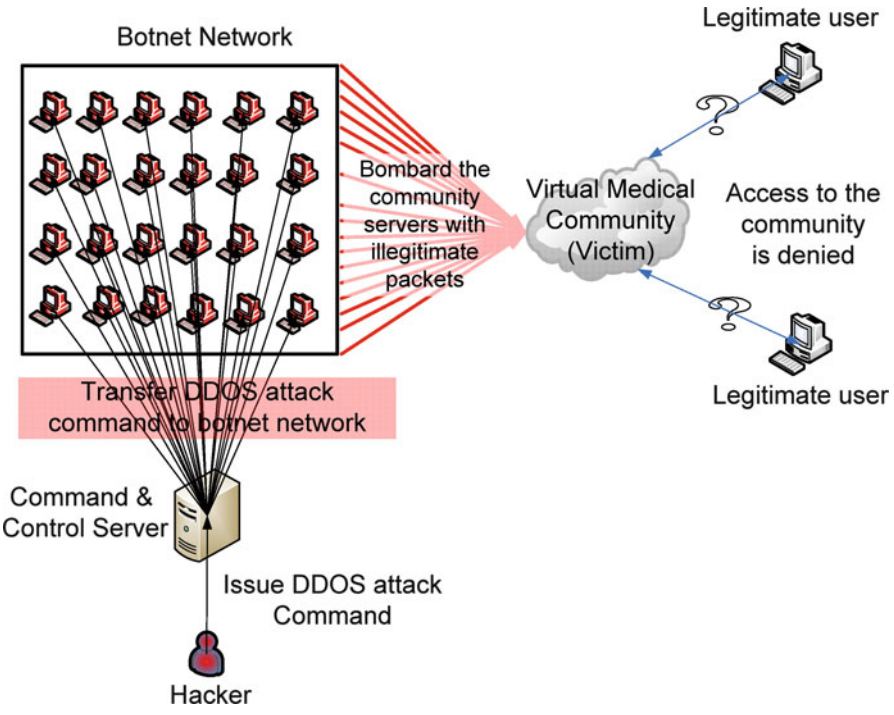

Fig. 5.8 A botnet attack against the virtual community

In such a scenario, as depicted in Fig. 5.8, the attacker logs into his Command and Control Center and transfers a distributed denial of service attack command to his botnet network. The zombie PCs of the botnet network bombard the community's information systems with illegitimate packets. Even if a single exploitable vulnerability exists, the attack will succeed, and legitimate users will not be able to access the community's services.

5.7.5.2 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 disaster recovery plan should always be ready. In the attack shown in Fig. 5.8, the attacking botnet network aims to bring down the users' points of entry and usage, namely, the authentication proxy, the web server, or the application server (designated targets depicted with red arrows in Fig. 5.9).

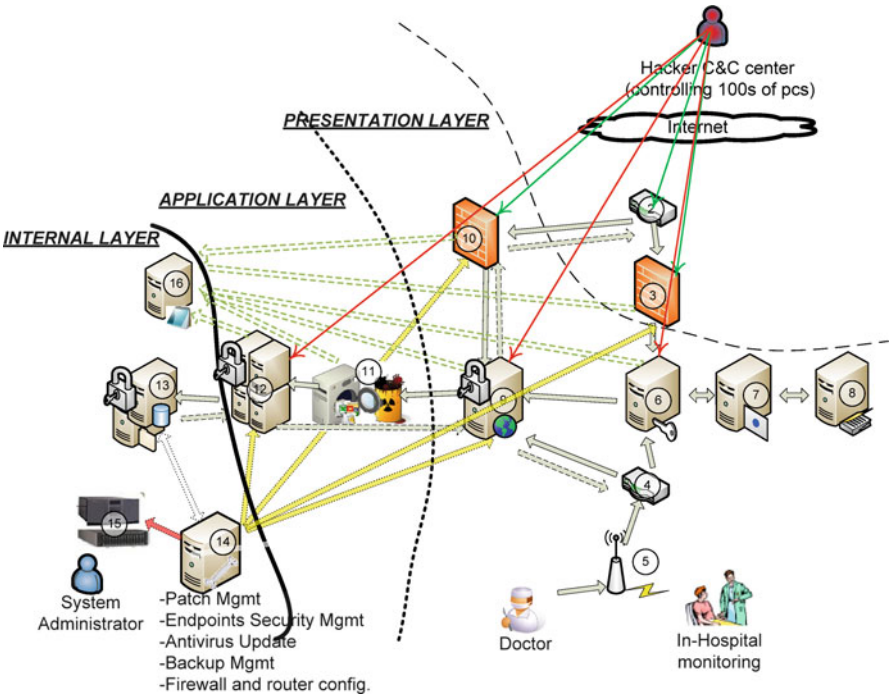

Fig. 5.9 The security infrastructure of the virtual community responds to a botnet attack

5.7.5.3 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 and use of BGP to block DOS attacks (see [41, 42]), 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 DDOS 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.

In the attack shown in Fig. 5.7, once the DDOS attack is launched, the community's defenses are activated, and the firewalls block the attack based on their proper configuration and their patching state (dashed green arrows in Fig. 5.9). This means that the employed patch management suite also plays a role (yellow arrows in

Fig. 5.9). If the attack is not blocked and the community goes down, the damage needs to be assessed and maintained, while the community must get back online (thick red arrow in Fig. 5.9). All these will be achieved based on the community's disaster recovery plan. The attack, successful or not, will be logged in the auditing server (dashed yellow arrows in Fig. 5.9).

5.7.5.4 Real-World Incident

In July 2008, an attack was started against several governmental, healthcare, and business web sites worldwide [9]. The attack was performed by bots infected by a malware and compromised by Asprox toolkit attack. The bots originated from over 1,000 unique web site domains including web sites of respectable organizations, governmental institutes, healthcare organizations (such as nhs.uk, samedical.org), etc. The attack toolkit is designed to inject a script into legitimate webpages. The malicious script exploits several vulnerabilities on the victim's machine in order to heighten the chances for successful exploitation: MDAC Vulnerability, QuickTime Rtsp Vulnerability, and AOL SuperBuddy ActiveX Control Code Execution Vulnerability. The successful execution of the script resulted in the downloading and execution of a Trojan Horse on the victim's machine.

The remedy to an attack of this type is to apply all the available patches in a daily basis and perform daily scans for malware and Trojans on all computers in the network. Patches may include configuration guidelines for firewalls and routers that may block illegitimate traffic.

5.7.6 Scenario 6: Malware

5.7.6.1 Description

This kind of incident scenario involves a computer system, whether a server system or a user PC, being infected by a malicious program (malware). The effects of any malware depend mainly on its payload which defines the type of infection (worm, Trojan Horse, Rootkit, spyware, adware, etc.) the infected system suffers from. The most dangerous infections are the ones the system cannot even see (Rootkits), the ones that come disguised as useful programs (Trojan Horses), and the ones that not only cause damage but spread also across the network to machines suffering of the same security hole (worms). A Trojan Horse can be used to steal medical data, a Rootkit can be used to control whole computer systems causing damage that cannot be calculated, while worms can cause a lot of damage by not only infecting a single computer system but also damaging similar systems effectively, wracking havoc among users, and disrupting entire systems with notorious examples: the Morris Worm, the Slammer Worm, etc. [50].

5.7.6.2 Critical Points

The backbone of the community's infrastructure (servers) as well as the end user's equipment might be the target of this attack. The attacker depending on the payload of the virus used will aim to exploit a single vulnerability preferably by using a zero-day exploit virus. The attacker could aim to steal information by means of a Trojan Horse or render the database server, for example, unavailable by using a worm, or control the RADIUS server by infecting it with a Rootkit.

5.7.6.3 Security Mechanisms Being Activated

The first step to guard against malware attacks is to have antivirus suites in place that protect both critical systems such as servers and end-users' systems, namely, members' systems. These antivirus suites must be kept up to date by frequently applying updates. This is performed, as designed, in the proposed security infrastructure for the community.

The second step is to have a properly configured patch management suite in place that operates in a regular basis, in order to limit system's vulnerabilities, which can be exploited by a malware through appropriate exploits. 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.

5.7.6.4 Real-World Incident

On May 2009, the Alberta Health Services (AHS) network in Edmonton has reported that a virus has infected 150 of their computers for a 2 weeks period.[9] The virus, which was previously unknown to the AHS antivirus software vendors, captured lab results, diagnostic imaging reports, and whatever else was on a computer screen and then transmitted the information to an external web site. Once the virus was detected, AHS worked quickly to remove the virus and take necessary steps to reinforce antivirus protection.

5.7.7 Scenario 7: Social Engineering

5.7.7.1 Description

In this incident scenario, a malicious outsider gains access to an organization's infrastructure by manipulating people working in the organization in a way that they happily reveal information which is otherwise confidential, such as Social Security

[9] DataLoss database report of incident 2174 [cited June 4, 2010]. Available from http://datalossdb.org/incidents/2174-personal-health-information-of-11–582-stolen-by-virus

Number and password. Such an attack could come from anywhere. For example, a person posing as a technician might call and persuade a user to reveal his password as part of a security maintenance procedure.

5.7.7.2 Critical Points

This kind of attack scenario aims mainly at the human factor. In order for a social engineer to be successful, he must be able to access physically or electronically a user of the community that has something useful to divulge, such as passwords and sensitive data. In the case of the proposed virtual community, targeted items could be means of authentication and data related to end users or doctors.

5.7.7.3 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 obtaining a password through a social engineering attack. The attacker would have to convince the user to hand on his e-token device as the log-in procedure requires something the user has (e-token) and something a user knows (password). The e-token device adds an extra authentication layer, as an attacker who learns something a user knows (password) is still denied access to the community's systems. Additionally, even if a social engineer manages to enter the site and tries to convince users to hand on sensitive data, 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 social engineer. The existence of an overall access and behavior policy educates the user and prevents him from being an easy prey to social engineering attacks. Finally, physical security measures deter any attacker from entering the community's IT room unaccompanied and gaining physical access to sensitive data.

5.7.7.4 Real-World Incident

On May 2009, a security consultant working in Siemens targeted a client that provides financial services [40]. His goal was to establish what level of access to information he could gain by means of social engineering. Without using any special equipment, he walked into the company's premises, established a base in a meeting room, and managed to access various company's premises, with most indicative the company's data room, IT, and telecoms network. He also managed to obtain usernames and passwords from 17 out of 20 company's employees by posing as a member of an IT department. He even managed to assess the working state of the company's CCTV circuit and bring a second consultant into the company's premises.

5.8 Discussion and Conclusions

This chapter presented a security-enabled architecture for a virtual healthcare community. The architecture is associated with a risk management model, which is based on the identification and protection of the community assets. In this community, patients have web access to the community services and provide their medical data using wireless sensor devices and/or web browsers, and doctors access community services either remotely or from inside the hospital and share their expertise with patients and other practitioners.

The design of a secure and trustful community is a difficult though interesting task, which should be preferably performed by following standardized procedures. In this direction, this chapter capitalizes on widely accepted security standards (the ISO 27000 family of standards) and provides a roadmap for developing a secure solution.

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. In this work, we presented the details of the application of the risk management model in a healthcare community and several security violation incidents, in healthcare and other sectors, which illustrate the various security mechanisms and validate our security model.

In addition to the security model, the risk management model, and the scenario-based evaluation, we presented a study on the legal implications of security violation incidents and introduced the process followed by the Hellenic Data Protection Authority for preventing and handling violators.

The next steps of this work comprise the prototype implementation of the security model for a healthcare community and an evaluation that will cover all possible attack scenarios.

References

1. Akyildiz, I. F., Su, W., Sankarasubramaniam, Y., & Cayirci, E. (2002). Wireless sensor networks: A survey. *Computer Networks, 38*(4), 393–422.
2. Apostolakis, I., Chryssanthou, A., & Varlamis, I. (2009). A holistic perspective of security in health related virtual communities. In A. Lazakidou & K. Siassiakos (Eds.), *Handbook of research on distributed medical informatics and E-health* (pp. 367–381). Hershey, PA: IGI Global.
3. Article 29 Data Protection Working Party. (2010, February). Opinion 1/2010 on the concepts of “controller” and “processor”. Brussels, Belgium.
4. Becker, M. Y., Fournet, C., & Gordon, A. D. (2007, July 6–8). Design and semantics of a decentralized authorization language. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF 07), Venice, Italy* (pp. 3–15). Washington, DC: IEEE Computer Society.

5. Becker, M. Y., & Sewell, P. (2004, June 28–30). Cassandra: Flexible trust management applied to electronic health records. In *Proceedings of the 17th IEEE Workshop on Computer Security Foundations; Asilomar Conference Center, CA, USA, 2004* (pp. 139–154). Washington, DC: IEEE.
6. Blaze, M., Kannan, S., Lee, I., Sokolsky, O., Smith, J. M., Keromytis, A. D., et al. (2009, February). Dynamic trust management. *IEEE Computer Magazine*, *42*(2), 44–52.
7. Chryssanthou, A., Latsiou, C., & Varlamis, I. (2009, June 9–13). Security and trust in virtual healthcare communities. In *Proceedings of the 2nd International Conference on Pervasive Technologies Related to Assistive Environments (PETRA 09), Corfu, Greece* (pp. 1–8). New York: ACM Press.
8. Curtis, D. W., Pino, E. J., Bailey, J. M., Shih, E. I., Waterman, J., Vinterbo, S. A., et al. (2008). SMART – An integrated, wireless system for monitoring unattended patients. *Journal of the American Medical Informatics Association, 15*(1), 44–53.
9. Cyberinsecure.com. (2008, July 18). *Asprox botnet mass attack hits governmental, healthcare, and top business websites* [cited 2010 June 4]. Retrieved July 30, 2010, from http://cyberinsecure.com/asprox-botnet-mass-attack-hits-governmental-healthcare-and-top-business-websites/
10. DataLoss Database. (2010, June). *Open security foundation; c2005–2010*. Available from http://datalossdb.org
11. Demiris, G. (2005). Virtual communities in health care. In B. Silverman, A. Jain, A. Ichalkaranje, & L. Jain (Eds.), *Intelligent paradigms for healthcare enterprises* (Germany-studies in fuzziness and soft computing, Vol. 184, pp. 121–137). Berlin/Heidelberg: Springer.
12. Demiris, G., Parker, O. D., Fleming, D., & Edison, K. (2004). Hospice staff attitudes towards telehospice. *The American Journal of Hospice & Palliative Care, 21*(5), 343–348.
13. Dixon P (2006, March 3). *Medical identity theft: The information crime that can kill you* [cited 2010 January 4]. The World Privacy Forum. First report in a series [Internet]. Cardiff by the Sea, CA, USA: World Privacy Forum. Retrieved May 22, 2011, from http://www.worldprivacyforum.org/pdf/wpf_medicalidtheft2006.pdf
14. Ebner, W., Leimeister, J. M., & Krcmar, H. (2004, January 5–8). Trust in virtual healthcare communities: Design and implementation of trust-enabling functionalities. In *Proceedings of the 37th Hawaii International Conference on System Sciences (HICSS 04) – Track 7, Big Island, Hawaii* (p. 70182). Washington, DC: IEEE.
15. European 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. *Official Journal of the European Communities, 38*(281), 31.
16. European Council. (1997). *Explanatory memorandum to recommendation (97) 5 on the protection of medical data*. Strasbourg, France: Council of Europe.
17. Federal Trade Commission. (2000). *Identity Theft Victim Assistance Workshop*. Washington, DC: Federal Trade Commission; c2000–2010 [cited 2010 June 4]. Retrieved June 4, 2011, from http://www.ftc.gov/bcp/workshops/idtheft/
18. Hitrustalliance.net. (2009). *Frisco, TX: Health Information Trust Alliance* [updated 2010; cited 2010 June 4]. Available from https://www.hitrustcentral.net/
19. ISO/CD. (2005). *ISO/CD 22857:2004: Health informatics – Guidelines on data protection to facilitate trans-border flows of personal health information*. Geneva, Switzerland: ISO/CD.
20. ISO/IEC. (2005). *ISO/IEC 27001:2005: Information technology – Security techniques – Information security management systems – Requirements*. Geneva, Switzerland: ISO/IEC.
21. ISO/IEC. (2005). *ISO/IEC 27002:2005: Information technology – Security techniques – Code of practice for information security management*. Geneva, Switzerland: ISO/IEC.
22. ISO/IEC. (2008). *ISO/IEC 27005:2008: Information technology – Security techniques – Information security risk management*. Geneva, Switzerland: ISO/IEC.
23. ISO/IEC. (2008). *ISO/IEC 27799:2008: Health informatics –Information security management in health using ISO/IEC 27002*. Geneva, Switzerland: ISO/IEC.

24. Jones, V. M., van Halteren, A. T., Dokovski, N. T., Koprinkov, G., Peuscher, J., Bults, R., et al. (2006). Mobihealth: Mobile services for health professionals. In R. S. H. Istepanian, S. Laxminarayan, & C. S. Pattichis (Eds.), *M-health emerging mobile health systems* (pp. 237–246). New York: Springer.
25. Kaplan, D. (2009, March 2). Group unveils first-of-its-kind standard to secure patient data. *SC Magazine, NEWS*. Retrieved March 4, 2009, from http://www.scmagazineus.com/group-unveils-first-of-its-kind-standard-to-secure-patient-data/article/128168/
26. Kui, M., Yue, W., Xu, Z., Xiaochun, X., & Gengdu, Z. (2005, September 21–23). A trust management model for virtual communities. In *Proceedings of the 5th International Conference on Computer and Information Technology (CIT 05), Shanghai, China* (pp. 741–745). Washington, DC: IEEE.
27. Kyriacou, E., Pavlopoulos, S., Berler, A., Neophytou, M., Bourka, A., & Georgoulas, A. (2003). Multipurpose health care telemedicine systems with mobile communication link support. *Biomedical Engineering Online, 2*, 7.
28. Laleci, G. B., Dogac, A., Olduz, M., Tasyurt, I., Yuksel, M., & Okcan, A. (2008). SAPHIRE: A multi-agent system for remote healthcare monitoring through computerized clinical guidelines. In R. Annicchiarico, U. Cortés, & C. Urdiales (Eds.), *Agent technology and e-health* (Whitestein series in software agent technologies and autonomic computing, pp. 25–44). Basel, Switzerland: Birkhäuser.
29. Law 2472/1997: Protection of individuals from personal data processing, Pub. L. No. 2472, Greece (1997).
30. Law 3418/2005: Medical code of deontology, Pub. L. No. 3418, Greece (2005).
31. Law 3471/2006: Protection of personal data and privacy in the telecommunications sector – Amendment of Law 2472/1997, Pub. L. No 3471, Greece (2006).
32. Lorincz, K., Malan, D. J., Fulford-Jones, T. R. F., Nawoj, A., Clavel, A., Shnayder, V., et al. (2004). Sensor networks for emergency response: Challenges and opportunities. *IEEE Pervasive Computing, 3*(4), 16–23.
33. Maji, A. K., Mukhoty, A., Majumdar, A. K., Mukhopadhyay, J., Sural, S., Paul, S., et al. (2008, January 29). Security analysis and implementation of web-based telemedicine services with a four-tier-architecture. *Proceedings of the 2nd International Workshop on Connectivity, Mobility and Patients' Comfort (CMPC), Tampere, Finland* (pp. 46–54). New York: ACM.
34. McClure, S., Scambray, J., & Kurtz, G. (2003). *Hacking exposed: Network security secrets and solutions* (4th ed.). Berkeley, CA: McGraw-Hill/Osborne.
35. Mondy, J., & Torresi, M. (2008). CIGNA creating a virtual health care community. *CIGNA website, News Releases*. Retrieved June 4, 2010, from http://newsroom.cigna.com/article_display.cfm?article_id=925
36. Mufti, M., Agouridis, D., Din, S., & Mukhtar, A. (2009, June 9–13). Ubiquitous wireless infrastructure for elderly care. In *Proceedings of the 2nd International Conference on Pervasive Technologies Related to Assistive Environments (PETRA 09), Corfu, Greece* (pp. 1–5). New York: ACM Press.
37. Ng, H. S., Sim, M. L., & Tan, C. M. (2006). Security issues of wireless sensor networks in healthcare applications. *BT Technology Journal, 24*(2), 138–144.
38. Orrin, S. (2004). *The twelve most common application level hack attacks*. Watchfire Corporation Whitepaper. Retrieved June 4, 2010, from http://www.emedia.co.uk/FM/GetFile.aspx?id=58740
39. Parducci, B., Lockhart, H., Levinson, R., & McRae, M. (2005). *eXtensible Access Control Markup Language (XACML) Version 2.0 core specification.* Billerica, MA: OASIS. Retrieved from www.oasis-open.org/committees/xacml/
40. Raywood, D. (2009, May 6). Social engineering attack allowed consultant to access company's data room and steal passwords. *SC Magazine, NEWS*. Retrieved June 4, 2010, from http://www.scmagazineuk.com/Social-engineering-attack-allowed-consultant-to-access-companys-data-room-and-steal-passwords/article/136278

41. *RFC2267 – Network ingress filtering. Defeating denial of service attacks which employ IP source address spoofing*. (2010). Available from Internet Engineering Task Force website. Retrieved January 2, 2004, from http://www.ietf.org/rfc/rfc2267.txt
42. *RFC3882 – Configuring BGP to block denial-of-service attacks*. (2010). Available from Internet Engineering Task Force website. Fremont, CA. Retrieved June 4, 2010, from http://www.ietf.org/rfc/rfc3882.txt
43. Schopp, L. H., Hales, J. W., Quetsch, J. L., Hauan, M. J., & Brown, G. D. (2004). Design of a peer-to-peer telerehabilitation model. *Telemedicine Journal and e-Health, 10*(2), 243–251.
44. Seamons, K., Winslett, M., Yu, T., Yu, L., & Jarvis, R. (2003). Protecting privacy during online trust negotiation. In R. Dingledine & P. Syverson (Eds.), *LNCS 2482: Proceedings of the 2nd Workshop on Privacy Enhancing Technologies (PET 2002), April 14–15, 2002, San Francisco, USA* (pp. 129–143). Berlin: Springer.
45. Stanberry, B. (1998). The legal and ethical aspects of telemedicine: Data protection, security and European law. *Journal of Telemedicine and Telecare, 4*(1), 18–24.
46. Stefanov, H., Bien, Z., & Won-Chul, B. (2004). The smart house for older persons and persons with physical disabilities: Structure, technology arrangements, and perspectives. *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 12*(2), 228–250.
47. U.S. Congress. (1996) Health Insurance Portability and Accountability Act, USA. Pub. L No. 104-191, 110 Stat. 1936.
48. US Department of Health and Human Services, Office for Civil Rights. (2003). *Standards for privacy of individually identifiable health information*. Washington, DC: US Department of Health and Human Services.
49. Varlamis, I., & Apostolakis, I. (2010). Self-supportive virtual communities. *International Journal on Web Based Communities, 6*(1), 43–61. doi:10.1504/IJWBC.2010.030016.
50. Vlachos, V., Spinellis, D., & Androutsellis-Theotokis, S. (2009, September 23–25). Biological aspects of computer virology. In A. P. Sideridis & C. Z. Patrikakis (Eds.), *Proceedings of the 3rd International Conference on e-Democracy, Athens, Greece* (pp. 202–219). Berlin: Springer.
51. Wang, X., Lao, G., DeMartini, T., Reddy, H., Nguyen, M., & Valenzuela, E. (2002, November 22). XrML – eXtensible rights markup language. In *Proceedings of ACM Workshop on XML Security (XMLSEC '02); Fairfax, VA* (pp. 71–79). New York: ACM.
52. Warren, S., Lebak, J., Yao, J., Creekmore, J., Milenkovic, A., & Jovanov, E. (2005, September 1–4). Interoperability and security in wireless body area network infrastructures. In *Proceedings of the 27th Annual International Conference of Engineering in Medicine and Biology Society (IEEE-EMBS), Shanghai, China*, *4,* 3837–3840.