Privacy-Preserving Mechanisms for Organizing Tasks in a Pervasive eHealth System*

Milica Milutinovic[1], Vincent Naessens[2], and Bart De Decker[1]

[1] KU Leuven, Dept. of Computer Science, DistriNet/SecAnon
firstname.lastname@cs.kuleuven.be
[2] Katholieke Hogeschool Sint-Lieven, Dept. of Industrial Engineering
firstname.lastname@kahosl.be

Abstract. In this paper, we describe privacy-preserving protocols for a scheduling service in eHealth applications. The scheduling mechanism that we propose protects sensitive information that are handled by the system. However, it still allows for a fair distribution of tasks and restricting the task assignment to caregivers with specific qualifications.

Keywords: eHealth, scheduling, privacy, fairness, caregiver, commercial.

1 Introduction

During the past decades, there has been a persistent trend in the age demographics of the developed countries. The average age of individuals is steadily increasing, and consequently there is a growing number of elderly that require continuous help. However, their guardians are often not in a position to provide this kind of daily assistance. Therefore, the potential of eHealth systems to provide efficient and cost-effective care in the home was recognized quite early.

In order to sustain certain independence of the elderly or patients, the home assistance systems should provide a wide range of services. These include continuous monitoring of health parameters, their automatic assessment and detection of anomalies, remote access to this data by authorized medical personnel, and communication with the caregivers. In addition to this, the patients or elderly need to be able to request assistance and these specific tasks should be assigned to their caregivers in a fair way. Examples are regular doctors' visits, catering or cleaning services, help with administering medication, etc. In order to take the burden of the close relatives of the patients, who are usually handling the organization of tasks, we explore its delegation to an eHealth system. Of course, for a pervasive system that will handle sensitive patient's data, such as health parameters or contacts with caregivers, one of the most important requirements

* This research is partially funded by the Interuniversity Attraction Poles Programme Belgian State, Belgian Science Policy, by the EU FP7 project NESSoS, by the Research Fund KU Leuven and by the IWT-SBO Project DiCoMas (Distributed Collaboration using Multi-agent System Architectures).

D. Gritzalis, S. Furnell, and M. Theoharidou (Eds.): SEC 2012, IFIP AICT 376, pp. 555–560, 2012.

is protection of the privacy. Therefore, we are proposing privacy-preserving protocols that describe the scheduling service in a pervasive eHealth system. Additionally, the described protocols surpass the need to disclose any identifying information to the scheduling service, which allows for it to be offered by a commercial entity. Nevertheless, the patients are still able to specify required skills and medical qualifications and their preferences regarding caregivers.

The **contribution** of this paper is the design of a few protocols that allow the integration of a scheduling service into a pervasive but privacy-friendly eHealth system. In current eHealth systems or research, privacy is either not an issue or is at best protected through access control mechanisms. However, accidental or malignant leakage of privacy-sensitive data (e.g. by hackers or employees) can be detrimental to the patient's privacy. In our system, we ensure that such leakage cannot occur. All the necessary information is available, but is only accessible to authorized caregivers that are directly involved with the patient's care.

Due to space limitation we will not discus related work. More extensive overview can be found in [2].

2 The System Architecture

An architecture for a pervasive eHealth system offering a wide range of services is proposed in [1]. Its high level description is given in this section and the development of the scheduling service is considered in this framework.

The proposed system design consists of four tiers. Patients' health parameters are continuously being measured by wearable, unobtrusive *sensors*. The recorded measurements are sent to and gathered by a *base station*, which represents a gateway towards the rest of the system. The base station can log and assess the measurements and request for a caregiver to be notified if a problem is detected. It also maintains patient-specific policies, which specify access control to stored data and thresholds and actions for normal, alert and emergency conditions.

The base station further communicates with a *dispatch centre* (DC), which can be a commercial organization that provides technical support and mediates the communication between the patients (i.e. the base stations) and the caregivers. It notifies the caregivers of tasks assigned to them or in case of emergency situations and follows up on their responses. It also relays the caregivers' requests to access the stored health-related data to the base stations. The data exchange is encrypted, in order to protect the information from the DC and its staff.

Finally, a separate entity, namely the *administration centre*, handles administrative tasks, such as user registration. This functionality is separated from the dispatch centre for privacy reasons. Since the dispatch centre mediates communication between patients and their caregivers, combining that information with the identities of all users would reveal sensitive data. For instance, knowing a specialist that is treating the patient reveals information about the patient's health. Therefore, all the users are identified in the dispatch centre only by their pseudonyms, but can be identified by a deanonymizing authority using the registration transcripts, when certain conditions are fulfilled.

Both dispatch and administration centres are equipped with tamper-free devices (TD_{DC} and TD_{AC}, respectively), which are used to re-key or decrypt sensitive information protected with their public key. These actions are possible only under certain conditions and are performed after strict checks in order to protect against attacks and prevent leakage of private information. For a detailed description of the functioning of these trusted devices, we refer the reader to [1].

2.1 System Functionality

Initially, all parties need to register with the administration centre. Upon registration, patients obtain a *smart card* recording their identity, address information and the service level agreement (SLA). The card also generates two public/private key pairs, one to be used for encryption, the other for signing. The public keys are then certified by the administration centre and the certificates are stored on the card. Similarly, the caregivers receive an *anonymous credential* certifying their identity, address information, medical qualifications and related proofs and a chosen random value which is not disclosed to the administration centre. The random value is used to generate provable unlinkable pseudonyms. Both patients and caregivers are then able to register with the dispatch centre.

When a patient registers with the dispatch centre (DC), her personal network is created and identified with her randomly-looking pseudonym. Linked with this pseudonym are a *vault*, i.e. an encryption of her real identity and address information obtainable only by the TD_{DC}, and a certificate issued by the TD_{DC}, linking the patient's pseudonym with her public keys. The trusted device only re-encrypts the information from the vault with the public key of an authorized party. The same holds for the caregivers' nodes. A caregiver anonymously approaches the dispatch centre when he is invited to join a patient's network. He then creates a new pseudonym as a function of the patient's pseudonym and the (secret) random number recorded in his credential. This way, he can pseudonymously authenticate with the DC, proving that the pseudonym is generated using his credential. Along with the caregiver's pseudonym, the DC stores his role[1] in the patient's network and a vault containing the identity, address, medical qualifications and appropriate proofs which only TD_{DC} can decrypt. Additionally, the caregiver generates a new key pair to be used for communication with the patient he is connected with and of which the public key is certified by the TD_{DC}.

The dispatch centre keeps the information about a patient's connections with her caregivers in the form of a network. These networks are patient-centric and both patients and caregivers are only identified with their pseudonyms.

3 The Scheduling Service

The scheduling service allows for fair allocation of patient-requested tasks to their caregivers, ensuring that every task is assigned and confirmed by a caregiver

[1] These roles are assigned by the patient at the time the connection is created. Examples are relative, neighbour, GP, specialist, cleaning service, etc.

with appropriate qualifications. It also allows the patients to specify the preferred caregivers or undesired ones. Therefore, the scheduling is offered by the system as a service the patients can subscribe to. The scheduling service can be offered by an external entity that registers with the dispatch centre. A patient can request to connect to it if she wants to use its services. In that case, the scheduling service becomes one of the caregivers of her network with an appropriate role. That is also true if the scheduling is offered by the dispatch centre itself.

3.1 Patient's Schedules and Caregiver's Profiles

The base station of every patient maintains a *schedule* of the caregivers' tasks. For every task, the schedule contains its identifier, the time frame, other details and the chosen caregiver. Until a caregiver confirms the assignment of a task, it remains conditional. For the task assignment, the base station may send a request to the scheduling service. However, the information about the tasks and the caregivers that is communicated to the scheduling service needs to be limited.

In order to allow the scheduling service to be offered by the system, caregivers specify their availability and store this information in a *profile* at the dispatch centre. They update this information regularly and the base stations or the scheduling service are able to retrieve the latest versions when a task is to be assigned. The profiles are linked with the caregiver's pseudonym and are stored encrypted with a secret key obtainable only by an authorized party (see further).

Besides the availability, the profiles contain other relevant information. They also state the caregivers' willingness to perform certain tasks and how long in advance they need to be notified about an assignment to be performed.

3.2 Protocol Assumptions

In the remainder of the text we will use the following assumptions about the creation of encrypted communication between two entities. All stakeholders in the system (patients and caregivers) interact via the dispatch centre (DC). This communication is SSL protected with server-side (DC) authentication. The initiating party (most often a caregiver) first pseudonymously authenticates with the DC. It then generates a fresh symmetric key and creates a *capsule*. The term 'capsule' denotes a symmetric key possibly concatenated with additional information, such as identifiers of the sending/receiving party, which is encrypted with the public key of the receiver. Next, the initiator's request (e.g. to access the patient's medical data) and the capsule are signed and sent to the DC. If the DC authorizes the request (based on the role of the sender), it relays the messages to the receiver. The receiving party then verifies the signature and extracts the session key from the capsule which is subsequently used for encrypting the communication between the two parties. We will refer to this encrypted communication between two parties as the creation of a *protected virtual link*. Once this link is established, all the messages DC relays are hidden from it.

A similar approach is used when some data needs to be stored at the DC in order to make it available to different parties. The entity that is sending the data

initially authenticates pseudonymously to the DC. It then generates a symmetric key with which it encrypts the data to be stored. The key is enclosed in a capsule along with the pseudonym of the sender and possibly the parties authorized to access the data. The capsule is encrypted with the public key of TD_{DC}. Both the encrypted data and the capsule are stored at the DC. The trusted device can make the protected data available to an authorized party by re-encrypting the capsule with this party's public key. However, this is only performed after thorough checks.

3.3 Scheduling a Task

If one or more tasks need to be assigned to the caregivers of a patient, the base station sends a request to the scheduling service via the dispatch centre. For every task, the base station creates task-assignment request containing the specified task, the time frame, the required caregiver's role and/or qualifications and preferred or undesired pseudonyms. This request is then signed and sent via a protected virtual link (see Sect. 3.2). It is then used by the scheduling service to prove its authorization to retrieve the profiles of the concerned caregivers. After verifying the request, signature and scheduling service's public key certificate, the trusted device re-encrypts the profiles' encryption keys for the scheduling service. If special qualifications are necessary for a task, the scheduling service can prompt TD_{DC} whether a particular caregiver has these qualifications. The TD_{DC} will check the caregiver's vault and reply 'Yes' or 'No'. When the scheduling provider obtains the profiles, it can assign the task, taking into account the required role of the caregiver and patient's preferences. Additionally, along with the initial request, the base station sends to the scheduling service relevant policies which are to be taken into account. Examples are limitation of hours that can be assigned to a role or a caregiver, restrictions on using commercial providers and additional requirements. In order to ensure a fair distribution of tasks, the base station also sends a summary (e.g. total number of hours) of current and past assignments for each of the caregivers. This way, the scheduling service can consider the load that is placed on each of the caregivers.

The assignments are then sent via a protected virtual link to the base station. The task assignments are now stored in the schedule, but remain conditional until the assigned caregiver approves the task.

3.4 Retrieval of Assignments

Every caregiver is assumed to be collecting his assignments at regular intervals (e.g. twice a week) via a web application. This dedicated application establishes a protected virtual link with the base station of the patient in order to retrieve the tasks for which the caregiver's pseudonym was assigned. The caregiver can then inspect them and reply, i.e. accept or reject the assignments. Subsequently, the web application needs to make changes to the availability information specified in his profile. For this, it authenticates pseudonymously to the DC, sends the new profile and corresponding capsule, both signed with the caregiver's anonymous

credential, and receives a receipt from the DC. This way, if a problem is detected, the DC can prove that the schedule was received from the caregiver, and the caregiver can also prove whether or not a version of the profile was sent by him.

However, when a caregiver wants to retrieve his complete profile, which might be the case if he is using the web application from a new device, strict checks are necessary. First, he will pseudonymously authenticate to the DC using his anonymous credential. The DC then asks the TD_{DC} to verify whether the pseudonym matches the one in the capsule. If this verification passes, the TD_{DC} re-encrypts the original capsule with the public key of the caregiver, which is sent to him with the profile. This way, only owners of the profiles are able to retrieve them.

3.5 Contacting a Caregiver

In case of emergency situations, a caregiver is urgently required to assist the patient. In this case, the base station chooses a caregiver itself, according to the patient's policies, so that no valuable time is lost. The base station sends an encrypted request to the dispatch centre to contact the chosen caregiver. The request can only be decrypted by TD_{DC}, which re-encrypts it for a calling module installed at the DC. For communication with the calling module, a symmetric key is used, which is embedded in the module through whitebox cryptography. The module sends the message specified in the request to the caregiver's phone number. This number is extracted by the TD_{DC} after appropriate checks, from the caregiver's vault stored at the DC. Once the message has been sent, the calling module will store the patient's pseudonym and a secure hash of the telephone number and delete any other information to protect from eavesdroppers. This hash is then compared to the incoming messages' origin. If the expected response is received, the calling module sends it encrypted to the base station.

For a detailed description of the protocols and evaluation of the security and privacy properties, we refer the reader to [2].

4 Conclusion

In this paper we have described a scheduling mechanism that can be integrated into a pervasive eHealth system. The focus of the design was preserving privacy of patients, but also their caregivers. Furthermore, the disclosure of information is performed on a need-to-know basis allowing the service to be offered by a commercial company, which is an important impetus for its deployment on a large scale.

References

1. Milutinovic, M., Decroix, K., Naessens, V., De Decker, B.: Commercially-run home assistance centres. Technical Report, vol. CW612, KU Leuven (2011)
2. Milutinovic, M., Naessens, V., De Decker, B.: Privacy-preserving scheduling mechanism for eHealth systems. Technical Report, vol. CW618, KU Leuven (2012)