NTNU
Innovation and Creativity

Developing Patient Controlled Access

An Access Control Model for Personal Health Records

Torstein Jensen
Knut Halvor Larsen

Master of Science in Computer Science
Submission date: June 2007
Supervisor: Øystein Nytrø, IDI
Co-supervisor: Gunnar René Øie, IDI

Norwegian University of Science and Technology
Department of Computer and Information Science

Problem Description

The personal health record is an evolving technology addressing problems with the handling and processing of information in the health care sector. Interest in this technology is growing, but the solution brings along many challenges that must be met. One challenge is the protection of patient data in shared systems and open networks, especially against threats that appear in such environments. The fact that the personal health record will have a diversity of users, including patients as well as healthcare personnel, also presents a challenge. The fact that the personal health record will have an increased diversity of users, as patients are included in addition to health care actors, also presents a challenge.

Another challenge is to adjust the PHR to local conditions. As the patients are given more control over access to their records, the personal health record requires modifications and adjustments according to local structure, settings and legislation. All these challenges must be met at the same time, and there are trade-offs between them.

This master thesis is a continuation of the project "Access Control Model for Personal Health Record" written autumn 2006. The research goals are as follows:

- Improve the tentative access control model from the previous project, and adjust it further to Norwegian conditions. Effective methods to reach this goal include interviews with patients.

- Determine whether the Indivo personally controlled patient record can be adapted to allow for the improved model, and if possible verify this with an implementation.

Assignment given: 22. January 2007
Supervisor: Øystein Nytrø, IDI

Abstract

The health and social care sector has a continuous growth in the use of information technology. With more and more information about the patient stored in different systems by different health care actors, information sharing is a key to better treatment. The introduction of the personal health record aims at making this treatment process easier. In addition to being able to share information to others, the patients can also take a more active part in their treatment by communicating with participants through the system. As the personal health record is owned and controlled by the patient with assistance from health care actors, one of the keys to success lies in how the patient can control the access to the record.

In this master's thesis we have developed an access control model for the personal health record in a Norwegian setting. The development is based on different studies of existing similar solutions and literature. Some of the topics we present are re-introduced from an earlier project. Interviews with potential users have also been a valuable and important source for ideas and inspiration, especially due to the fact that the access control model sets high demands on user-friendliness. As part of the access control model we have also suggested a set of key roles for the personal health record.

Through a conceptual implementation we have further shown that the access control model can be implemented. Three different solutions that show the conceptual implementation in the Indivo personal health record have been suggested, using the Extensible Access Control Markup Language as the foundation.

Preface

This thesis is written for the subject TDT4900 Computer and Information Science, Master's Thesis, at the Norwegian University of Science and Technology (NTNU), spring 2007. The thesis carries the study credit of one semester, and is the final subject of the MSc. level degree Master i teknologi / Sivilingeniør.

The assignment is given by the Department of Computer and Information Science (IDI,NTNU), and has been defined together with Research Fellow Gunnar René Øie at IDI. Øie has also been our supervisor throughout this thesis.

We wish to thank the following for inspiration, comments, rich discussions and information:

- Gunnar René Øie, Research Fellow, IDI, our supervisor
- Jorunn Bjerkan, Research Fellow, NSEP, cooperation and help with interviews
- Amund Tosterud and Marie Richter, help with interviews
- All the participants in the qualitative interviews

Trondheim,
19. June 2007

.....	.....
Knut Halvor Larsen	Torstein Jensen

Contents

List of Figures

List of Tables

CHAPTER 1

Introduction

This chapter presents the context, motivation and background for our master's thesis. It further presents the assignment and the report outline.

1.1 Context, Motivation and Background

The continuous progress within the health and social care sector has put larger and larger demands to medical knowledge, treatment and efficiency. This has among other things resulted in the introduction of computerised aids like the electronic health record (EHR). Helping and improving in certain areas, this introduction has also created new and important problems to address.

Some of the current EHRs, used by all Norwegian hospitals, have received criticism for their lack of privacy protection. Sensitive and personal information about a patient is often available for health care actors without any medical relationship to the patient. The systems are also developed based on the need for information availability for health care actors, which lead to a too open and wide access. On the other hand, legislation and lack of interoperability between systems restrain the information flow among health care actors. And while helping the health care actors in their tasks and efficiency, the patient is still thought of as a passive participant in the treatment process.

With the aim to give the patient a more central role in treatment, and making a system not only for the health care actors, the idea of the personal health record (PHR) has been introduced. By letting the patients have access and control of their own medical record, they can take a more active part in their own treatment, and at the same time adjust access to their medical information.

The objective of this master's thesis is to develop an access control model for the PHR from a Norwegian perspective. The development has a special focus on the users and their opinions,

and qualitative interviews will be the main source for this information. The thesis is a continuation of our work in the autumn project "Access Control Model for Personal Health Record" [1].

1.2 The Assignment and Research Goals

The following is the assignment text and research goals:

> The personal health record is an evolving technology addressing problems with the handling and processing of information in the health care sector. Interest in this technology is growing, but the solution brings along many challenges that must be met. One challenge is the protection of patient data in shared systems and open networks, especially against threats that appear in such environments. The fact that the personal health record will have a diversity of users, including patients as well as healthcare personnel, also presents a challenge. The fact that the personal health record will have an increased diversity of users, as patients are included in addition to health care actors, also presents a challenge.
>
> Another challenge is to adjust the PHR to local conditions. As the patients are given more control over access to their records, the personal health record requires modifications and adjustments according to local structure, settings and legislation. All these challenges must be met at the same time, and there are trade-offs between them.
>
> This master thesis is a continuation of the project "Access Control Model for Personal Health Record" written autumn 2006. The research goals are as follows:
>
> - Improve the tentative access control model from the previous project, and adjust it further to Norwegian conditions. Effective methods to reach this goal include interviews with patients.
> - Determine whether the Indivo personally controlled patient record can be adapted to allow for the improved model, and if possible verify this with an implementation.

Through these two research goals we will continue our work with access control in the PHR. In the first phase we will continue the development of the access control model, making the necessary changes based on qualitative interviews with patients, health care actors and other potential users. We have decided to focus our work on the access control model itself, together with a limited set of roles, thus leaving our work on the specific resource and institution lists, and the access matrix at status quo.

In the second phase we will try to adapt the Indivo personal health record to the access control model. This work will consist of studies of the existing access control solution in Indivo, in addition to the technologies used. Due to the fact that Indivo is still at a beta state, the implementation will most likely be of a conceptual character.

1.3 Thesis Outline

Chapter 2 gives an introduction to the personal health record. It also presents two ongoing projects that have similarities to the personal health record, namely the Kjernejournal and Unique Sampro. In Chapter 3 we introduce the open source project Indivo, which is an implementation of the personal health record. The chapter also gives an introduction to the Extensible Access Control Markup Language and access control in Indivo. Chapter 4 presents some of our previous work related to this thesis. Preparations in relation to new interviews are described in Chapter 5, with the following analysis of these interviews in Chapter 6. Chapter 7 then presents the finalised access control model for the personal health record, while a conceptual implementation of our model in the Indivo system is presented and explained in Chapter 8. The interviews, the access control model, and the implementation is discussed in Chapter 9, along with suggestions for further work. We sum up our thesis with the final conclusion in Chapter 10.

CHAPTER 2

The Personal Health Record

In this chapter we introduce the personal health record concept, and present motivations for using such a system.[1] Further we look into work related to the PHR in Norway, especially the Kjernejournal and the SamPro project, which is based on the Individual Plan.

2.1 Introduction to the Personal Health Record

In today's modern health care information is a key factor to provide efficient and good care for patients. The problem is that by law information about a patient can not be kept in a shared system, and a patient's record is therefore scattered among all these. This means a lot of duplicated information and unnecessary work for both the patient and the health care actors.

The patient has the right to access his or her medical information stored at any health care actor. Because of today's cumbersome method to retrieve this information few patients take an active part in their own record.

The introduction of a PHR aims to change these two problems. The PHR is a patient-centred system, where the record is owned and administrated by the patient. The record is not limited to a single organisation or a single health care actor, and can therefore receive and contain all the medical documents for the patient. The PHR will allow people to "access and coordinate their lifelong health information and make appropriate parts of it available to those who need it" [2]. The Markle Foundation's Connecting for Health collaborative has defined the PHR as: "An electronic application through which individuals can access, manage and share their health information, and that of others for whom they are authorised, in a private, secure, and confidential environment" [2]. The PHR is an area of large interest, and in USA the U.S.

[1] Section 2.1 and 2.2 are taken from our autumn project "Access Control Model for Personal Health Record" [1]

Secretary of Health and Human Services, the National Coordinator for Health Information Technology, and the Administrator of the Centres for Medicare and Medicaid Services have all identified PHRs as priority [3].

In addition to collecting relevant medical information in one place, the PHR gives the opportunity for the patient to contribute to the record. Information from the patient such as symptoms, medication, observations and actions can give the health care actor important and relevant information regarding diagnosis and treatment plans. This again will give a more active patient involvement with "better knowledge and ability to reflect upon their own health situation and medical problems and thus acts to a lesser degree as mere passive receivers of therapy" [4]. This way the patient does not only provide information, but also gains access to medical documents. The result could lower the barriers for communication between patients and health care actors and therefore make it easier to set up appointments, to request refills and referrals, and to report problems [3]. A prerequisite for this close cooperation is that medical documents are made understandable for the patient, something that will require some extra effort from the health care actors. Figure 2.1 illustrates the communication between the different actors.

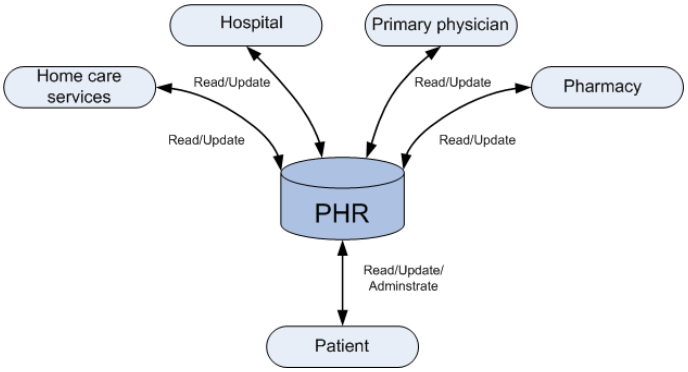

Figure 2.1: The personal health record

A PHR must have clear guidelines of who can do what. Figure 2.2 depicts this. Only health care actors have the ability to add or alter medical documents, while patient-provided information solely is in the hands of the patient. This creates a clear limitation of actions patients and health care actors can perform on documents in the PHR, illustrated by the horizontal line in the figure. The line represents a border where information exchange can only be done in a read only manner [4].

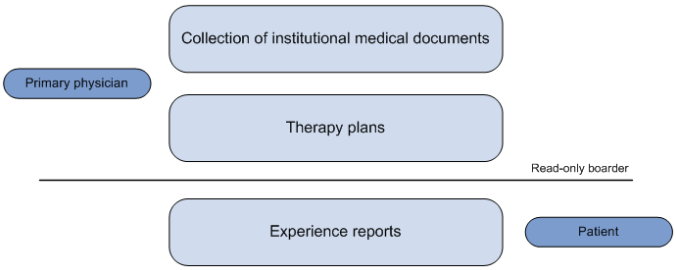

Figure 2.2: Ownership of information. Adapted from [4]

2.2 The Personal Health Record in Norway

Norwegian health care informatics has the last few years been driven by a sequence of strategic plans given from the Norwegian directorate for health and social affairs. Currently the use of EHR is widely spread throughout the sector as documentation of patient encounters is required. However, there has been a low level of electronic exchange of medical documents within the sector. Continuing strategic plans nationally and locally, and government-owned public companies such as the "Norwegian Centre for Informatics in Health and Social Care" (KITH), are now targeting this with the development of standardised messages and guidelines for electronic communication as well as smaller pilot projects. One of these pilot projects is the "Kjernejournal" which aims to simplify exchange of data between medical institutions. A focus on more patient centred care and the need to prepare individual plans for patients has resulted in the SamPro project, a co-operative project between SINTEF, Central Norway Regional Health Authority and Visma.

2.2.1 Kjernejournal

The Kjernejournal [5] is a solution that simplifies and speeds up communication between health care actors. In the future it might be a part of the EHR, and it also has some similarities with the PHR. In this section the Kjernejournal and its relation with the EHR is described. A comparison of the Kjernejournal and the PHR can be found in Section 2.2.2.

Background

Today the EHR in Norway is maintained by the patient's primary physician, and any new EHR relevant information is sent to him or her. This is done electronically, on paper or verbally. In many cases this information (e.g. a discharge summary) is sent after quite some time, and the quality of the content is not good enough. This leads to a situation where the primary physician does not have an up-to-date overview of the patient's medical history. In addition, he or she has to enter the information into the EHR, which often is too time consuming during a busy workday. If some other health care actor needs access to information about the patient from the EHR, the primary physician has to be contacted and together with the patient decide whether the request should be granted or not. An illustration of this situation is given in Figure 2.3.

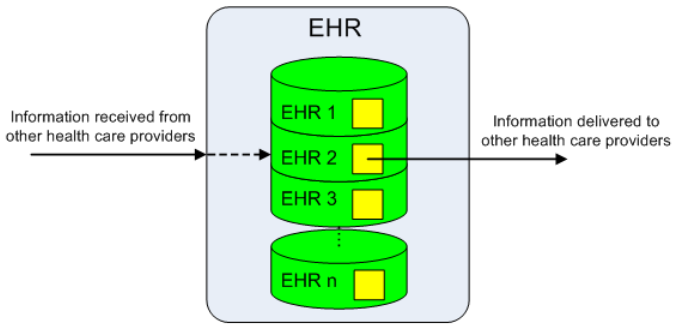

Figure 2.3: EHR today. Adapted from [5]

The Kjernejournal is a solution that addresses these issues. It is meant to simplify the message exchange between health care actors, and hence improve the quality of care by offering correct and up-to-date information to the requesting health care actors who treat the patient.

In Trondheim there is an ongoing project called "Fyrtårn Trondheim" where a Kjernejournal is being implemented. This is a simplified version that addresses the issue of adverse effects by sharing medicament related information [5].

Technical Solution

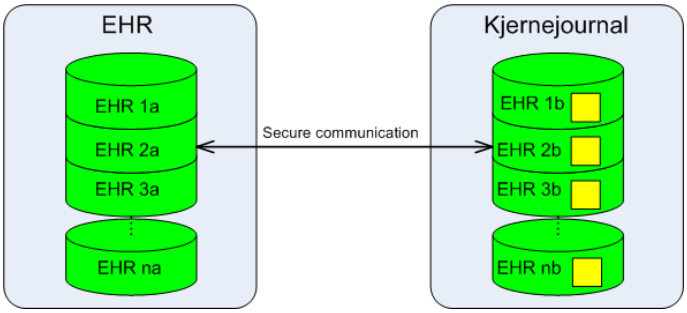

Figure 2.4: EHR and Kjernejournal. Adapted from [5]

Figure 2.4 illustrates how the Kjernejournal should be implemented. Information from the EHR that is relevant for other health care actors is copied from the EHR and placed on a central server that is physically separated from the original one. The Kjernejournal is, however, still a logical part of the EHR, which simplifies the handling of the Norwegian legislation. Communication between the two servers occurs by means of secure communication. To find out what information should be distributed, the patient sits down together with the primary physician. Together they identify a small number of situations where health care actors need access to information in the EHR. Some examples could be hospitalisation of the patient, and administration of medication. Then they decide what information should be shared for each situation and to whom. The Kjernejournal server must provide high availability in order to satisfy fast and reliable sharing of information.

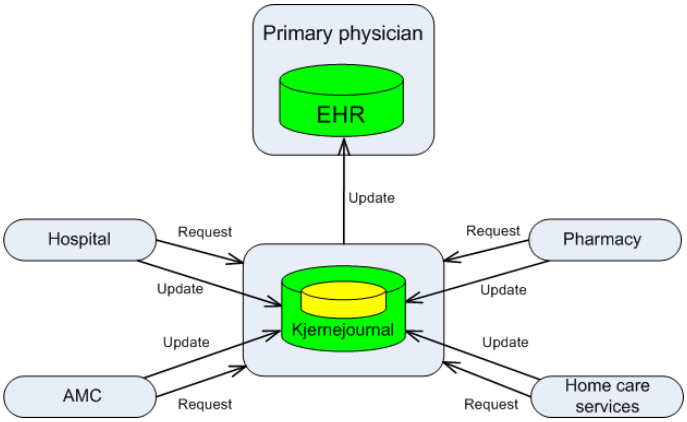

Figure 2.5: Communication with health care actors. Adapted from [5].

Figure 2.5 shows the communication between the different health care actor. When a request for information is made, the predefined rules are checked, and if the health care actor satisfies the requirements determined by the primary physician and the patient, the information is handed out. Hence the process is fully automated and there is no time-consuming communication with the primary physician. Access to the Kjernejournal could happen via the health care actor's traditional EHR or in some cases through the Internet.

When a health care actor adds a new document to the patient's Kjernejournal, the document is given a "not yet approved" status and a notification is sent to the primary physician. It is up to him or her to check the document and approve or reject it.

2.2.2 A Norwegian Approach to PHR

In Norway there is a growing interest in the use of PHR, and there is ongoing research on the subject at the Department of Computer and Information Science at the Norwegian University of Science and Technology (NTNU), in cooperation with The Norwegian EHR Research Centre (NSEP).

The Kjernejournal is a step in the direction towards a PHR. It is a tool where relevant medical information can be collected and distributed to health care actors who need this information. However, the Kjernejournal has some key shortcomings when it comes to the principles of a PHR. Brasethvik et al. point out three aspects [4]:

- The Kjernejournal does not have a system that allows patients access to the stored information. It is a mere collection of documents intended for exchange between health care actors.
- Since the patient does not have access to the Kjernejournal, there are no possibilities for patient provided information, one of the key features in a PHR.
- Documents in the Kjernejournal are not necessarily understandable for the patient. Medical documents are often written in a professional medical jargon, consisting of technical terms, abbreviations and incomplete sentences.

The Kjernejournal in its original form consequently does not support the patient or enhance the patient-physician communication as a PHR should do. However, since the Kjernejournal is under development in Norway, it is an excellent starting point for a Norwegian version of a PHR system. Brasethvik et al. suggest that a PHR called "EigenJournal" can be a supplement to the Kjernejournal to achieve the desired level of patient involvement. Figure 2.6 illustrates how this can be realised. There is a clear information ownership boundary between the Kjernejournal and the EigenJournal, but at the same time an information exchange between the two systems is allowed.

There are some problems combined with the use of the PHR as described above. Initially the interaction between the Kjernejournal and the PHR has to be clarified, especially when it comes to setting access to documents. Questions like "Should access control for the Kjernejournal and the PHR be the same?" and "Will changes made to one of them affect both?" will have to be answered.

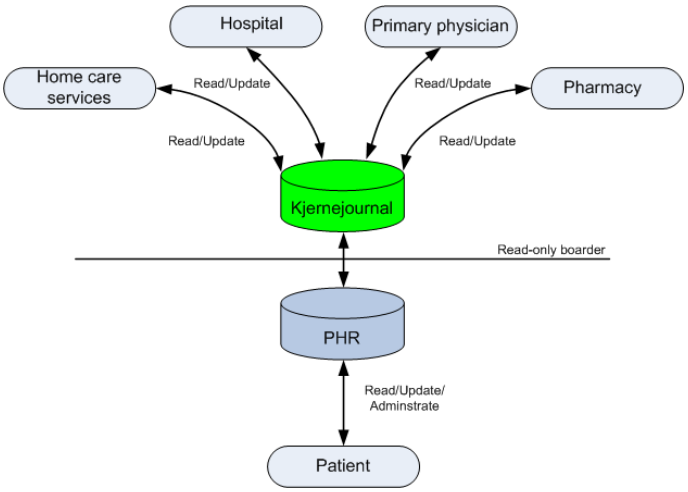

Figure 2.6: Kjernejournal and the PHR. Adapted from [4]

2.3 Unique SamPro

Unique SamPro is an electronic solution based on the Individual Plan (IP), with similar functionality as the PHR. In this section an introduction to IP is given, before introducing SamPro and it similarities and differences with the PHR.

2.3.1 The Individual Plan

The IP has in Norway been introduced to improve the conditions for patients in need of long-lasting treatment. This section explains what an IP is and why it has been necessary to introduce it. A description of the content in an IP document is also given.

Background and Legislation

Investigations of long-lasting treatment regarding rehabilitation and psychiatric suffering in Norway during the late 90s have revealed that the health care service is too divided and poorly coordinated. This leads to a situation where the receiver of health care services or his or her relatives have to use a lot of time in order to get the necessary help and treatment. They also experience that their qualifications are not taken into consideration when different solutions are discussed [6]. The IP is meant to address these issues. Its objectives include better adoption, coordination and evaluation of the provided services [7]. These objectives are realised by making the patient an active part in the decision making process around his or her treatment, as well as keeping regular plan meetings together with the involved service providers. One person from the plan is in charge of follow-up of both the service receiver and the plan. All the information in the plan is kept in a plan document, as described more in detail in the next section.

The IP Document

The IP document is not a goal in itself, but a tool needed in order to be able to create an adapted service offer to the service receiver [6]. It can be paper based or electronic, and it should not be too detail oriented. The document must contain several points. First of all, a description of the situation is required, with an overview of the the patient's goals, resources and needs for services. It is a good idea to let the goals be formulated by the service receiver, as he or she will feel greater ownership to the plan and thus be more motivated. The document must also contain an overview of participants in the plan and responsibilities. In addition to the service receiver and the service providers, it might in some cases also be natural to involve relatives. Especially if the relatives live together with the service receiver, or if the service receiver is a person under age.

Each IP has to have a person who is in charge of following up the service receiver and the plan. This person must be specified in the document. In most cases this will be the plan coordinator. Which services that are to be provided, their scope and demands regarding content, and the service provider in charge of each service must also be written down. The IP also has a start date and a time limit which must be specified. The service receiver has to give his or her consent that an IP is wanted and that the service providers may have access to the confidential information in the document. Finally, the form of co-operation between the different entities that are chosen to perform the services, must be described.

2.3.2 SamPro Background and Functionality

The SamPro project originally started as a co-operative project between SINTEF, Central Norway Regional Health Authority and the IT company Visma. The goal of the project was to develop a solution for individual care plans that could satisfy the legislation. Results of the project can today be seen in Visma's product Unique SamPro [8].

Unique SamPro is a web-based tool for patients under treatment and convalescence. It supports all the necessary requirements that a well written paper based IP does. However, Unique SamPro has additional functionality to further improve the work conditions for all actors involved [9]. This includes an access control system that enables the patient to share information with whomever he or she wants. To improve communication and co-operation between the actors, a message system has also been introduced. It has similarities with e-mail, but it is an internal system were all messages are stored on the Unique SamPro server. However, the sender of the message can choose to notify the receivers that a new message has arrived through Short Message Service (SMS) or e-mail.

In order to ease the coordination of the work, a calendar is provided. And as a result of requests from users, the newest version of Unique SamPro offers the possibility to write notes and blogs. The former feature can be used for minutes or notes from IP meetings, while the latter can be seen as the patient's diary. Finally, Unique SamPro logs all activity in the plan, and this can be seen in the log section. Unique SamPro is today actively used in the Central Norway Regional Health Authority. The main page is shown in Figure 2.7.

Although it is easy to access and use, Unique SamPro is not a complete patient record system, as will be explained in the following section.

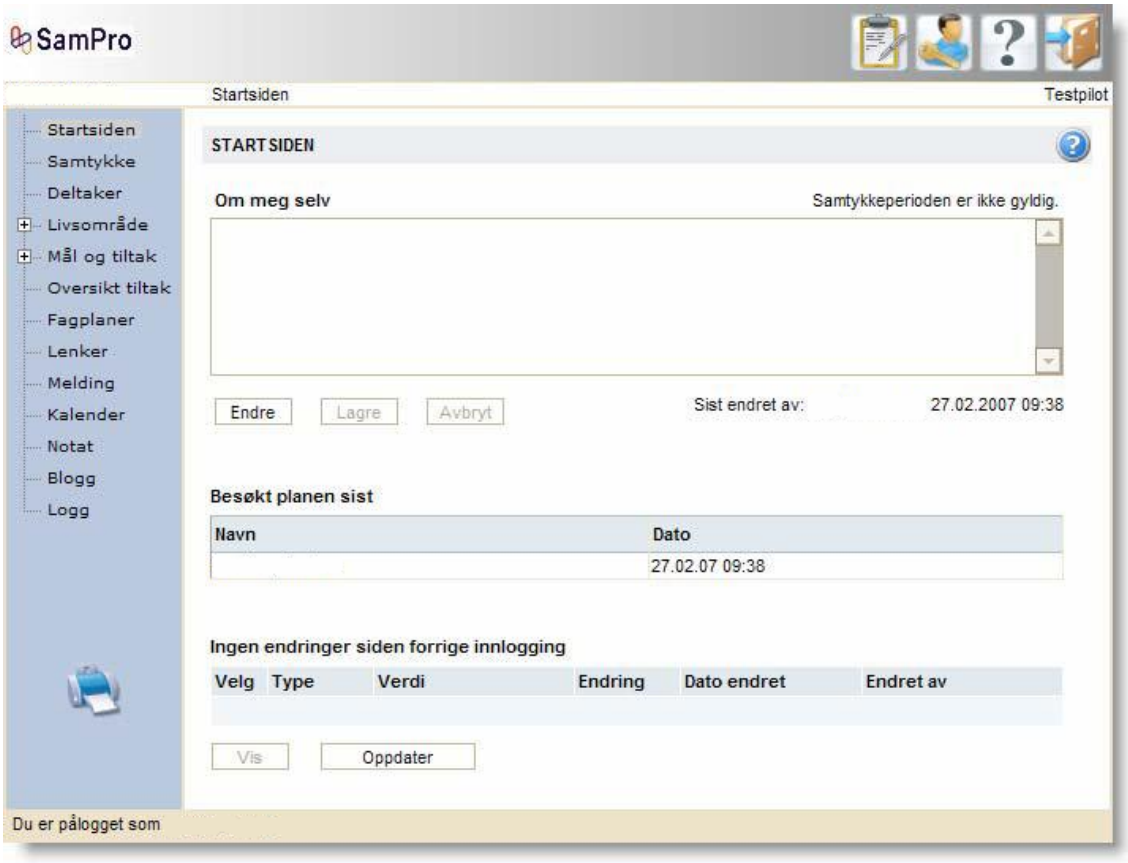

Figure 2.7: Sampro main page. Taken from [9]

2.3.3 SamPro and the PHR

Unique SamPro and the PHR have several similarities. They are both electronic, and can be accessed at all time by the actors involved. The patient has an active part in the system and can contribute with his or her own feedback and information. Further, the patient and the service providers can communicate with each other, and by using the access control system the patient can limit access to parts of the IP if necessary.

However, a PHR is supposed to contain all medical information about the patient. This is the main point that separates Unique SamPro from the PHR, since Unique SamPro only contains information relevant for the IP. In addition, some of the IP relevant documents are kept outside of SamPro and only referred to with description and physical location.

CHAPTER 3

Indivo

This chapter presents the PHR implementation Indivo with its architecture and access control, including the use of the access control policy language XACML [1].

3.1 Introduction

There are a few implementations of the Personal Health Record known today, and mainly these are at an early test stage [4]. One of these implementations is Indivo, which is a project at Harvard Medical School, MIT, and Children's Hospital Boston. In this section an introduction to Indivo is given with an overview of the system and an explanation of the technical architecture.

Indivo is basically designed for patients to have control over a complete secure copy of their medical record [10]. Hence it is not meant to be the primary record of the health care system but a collection of medical data across the patients' history. The patient is allowed to read, write and modify components, and he or she decides who shall have access to the different parts. This is done by granting rights to already registered persons, groups or roles. The access control can be done on a fine-grained level, for instance for each document in the patients medical record. Indivo is open source and free, and built to public standards. It is also module based and easy to configure to different needs. Indivo is currently under development, and the most recent version is 3.0 beta.

One of the main ideas with the PHR is to enable better communication between the primary physician and patient. Indivo supports this by including support for messages. It also supports a lot of other document types and new ones can easily be added due to Indivo's flexible structure. New information can be sent to the patient from any client (e.g. EHR system) as long as the information is sent using Indivo's communication protocol.

[1] Section 3.1 and 3.2 are adapted from our autumn project "Access Control Model for Personal Health Record" [1]

3.2 Indivo Architecture

Indivo can be divided into three layers: Client, server and store. Each layer can be located at different physical locations and both the client and the store are designed in a pluggable fashion such that a number of different types can be used. The ”heart and soul” of the architecture is the server since most of the data processing happens there.

Indivo uses a communication protocol called IndivoTalk to communicate between the client and the server. This protocol is based on Extensible Markup Language (XML) messages and request-response interaction. It can quite easily be extended by adding new XML schemas and performing small modifications to the server source code.

The following sections describe each layer of the Indivo architecture in more detail. Figure 3.1 serves as an illustration of the different layers and the interaction between them.

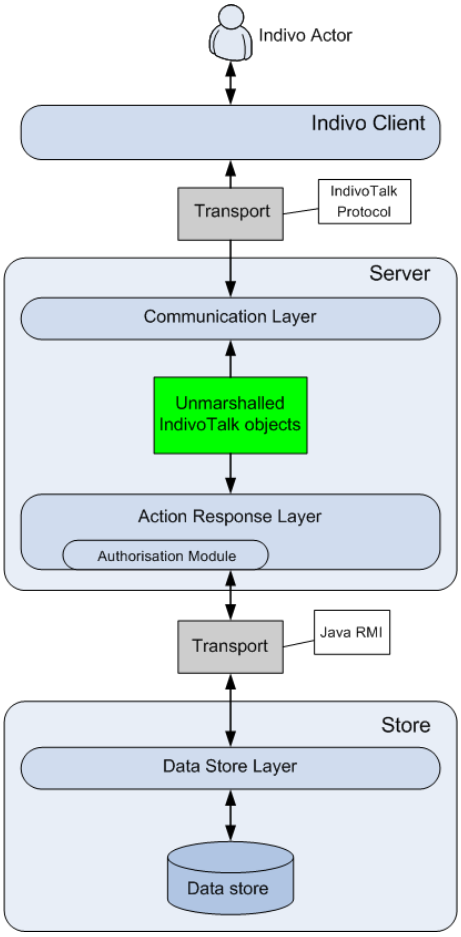

Figure 3.1: Indivo architecture

3.2.1 Client

The client is the layer used to communicate with the Indivo server. As long as it uses the IndivoTalk protocol the client can be any software process regardless of the platform and

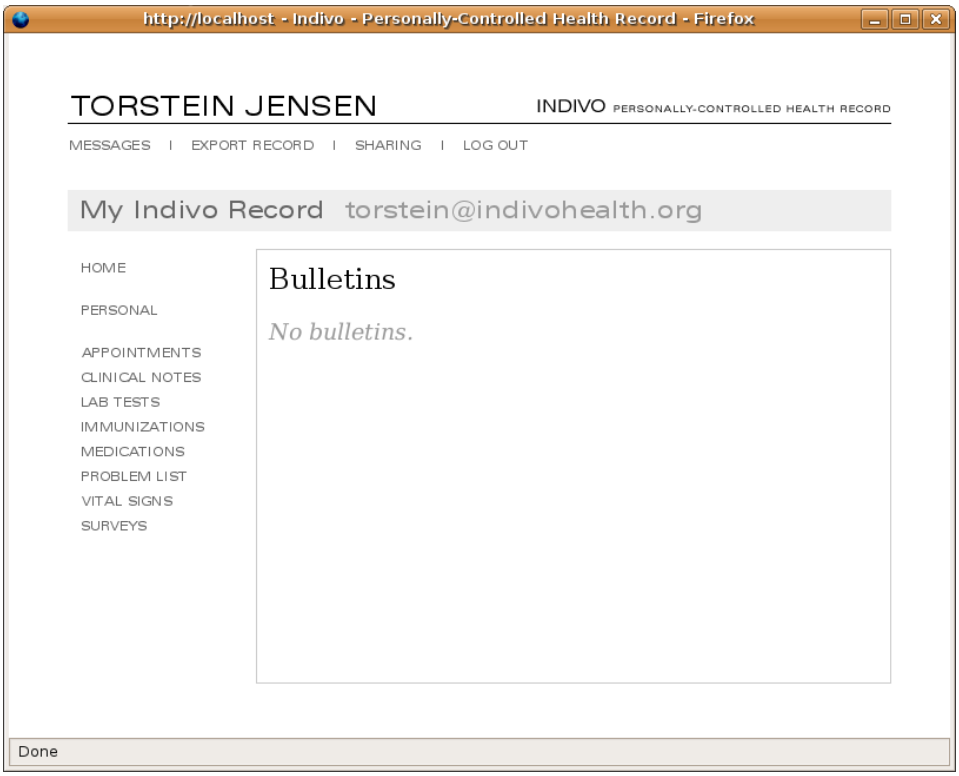

Figure 3.2: Indivo client

implementation language. A screenshot of the most recent client provided by the Indivo team is shown in Figure 3.2. This client is based on the Hypertext Preprocessor (PHP) language, and has limited functionality compared to what the server can offer. Among other things it does not support the fine-grained access control.

The Indivo actor is a registered user of the system. Each user has his own record and some attributes, e.g. roles and group memberships. Each role has some privileges. E.g. a user logging in with a patient role might be allowed to update, read and add documents to his/her own record, while a user logging in with a provider role must be given privileges by the patient in order to do the same to the patient's record.

3.2.2 Server

The server is a Java 2 Enterprise Edition (J2EE)-compliant Servlet. Using the servlet technology has several advantages. The Secure Hyper Text Transfer Protocol (HTTPS) can be used for encryption of IndivoTalk messages between the client and the server. Also the manipulation of requests is simplified by using the Java Servlet Application Programming Interface (API).

The server consists of the three different main parts: The Communication Layer, the Action Response Layer and the Authorisation Module.

The Communication Layer is responsible for accepting XML IndivoTalk messages and converting them into program objects which are then sent on to the Action Response Layer. When a response arrives from the Action Response Layer it is converted back into XML messages and

sent to the client. This processing is done using Java Architecture for XML Binding (JAXB). The messages are automatically marshalled and unmarshalled according to a specified schema which in this case consists of the different IndivoTalk elements.

The Action Response Layer processes the actions received from the Communication Layer and maintains information about the different sessions. There is one action responder for each type of action, and each one of these has its own processing and authorisation procedure. The layer reads several XML configuration files in order to create the different responders, and these files contain among other things information about the data store and authorisation engines. It is the responsibility of the Action Response Layer to delegate authorisation to the Authorisation Module before the request may be executed.

The Authorisation Module is an implementation of XACML, described in more detail in Section 3.3. Depending on the action responder (and type of action) there are either one or two authorisation steps. The first one decides whether the user is allowed to perform the type of action, and the second finds out whether there are any record policies that prohibit the user from performing the action. In the first case the authorisation policies are read from a configuration file, while the record based policies are read from the data store. The advantage of using XACML is that it is very flexible. A more complementary explanation of Indivo's access control can be found in Section 3.4.

3.2.3 Store

The data store contains the different records. Each record consists of several documents and each document consists of several versions. The different document types are defined with XML schemas at the server, and adding new document types is easy. At compilation time they are translated into Java classes using JAXB.

Technically speaking the default store is a Berkeley Database consisting of several encrypted XML files. They are accessible to the Action Response Layer using Java Remote Method Invocation (RMI). Encryption can occur before the data is sent from the server or at the store. One advantage of encrypting the data at the server is that the encryption key is kept separate from the encrypted data.

The default store can be replaced by other store types, as long as they adhere to the Indivo API which consists of method calls that the server uses when it communicates with the store.

3.3 XACML

This section gives a basic introduction to the Extensible Access Control Markup Language (XACML) defined in the OASIS standard eXtensible Access Control Markup Language (XACML) Version 2.0 [11].

3.3.1 Key Concept

XACML is a language for access control based on Extensible Markup Language (XML) [12] [13]. It defines both a policy language and an access control decision request and response

language. The policy language allows for simple and complex policies to be defined, supported by functions, data types and combining logic. A runtime query, represented in the request language, can then be compared to the relevant policy to either allow or deny specific actions based on functions comparing attributes. Figure 3.3 shows the scope of the XACML language indicated by the shaded area.

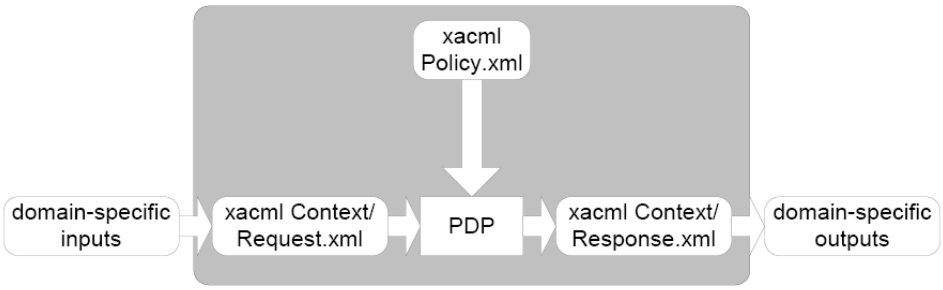

Figure 3.3: XACML context. Taken from [12]

A typical scenario in access control handling is that someone wants to take some action on a resource. This request is received by a system, e.g. a file system, which is defined as a Policy Enforcement Point (PEP). See the data-flow diagram in Figure 3.4. The PEP formulates a request based on the original requester's contributed information, called attributes, such as the resource in question and the action the requester wants to perform (1). Extraction of these attributes is done by the Policy Information Point (2). The request is then sent from the PEP to the Policy Decision Point (PDP), which evaluates the request (3). The PDP uses the policies and rules defined in the policy language (4) that applies to the request, and comes up with an answer about whether access should be granted or not(5). The final answer is returned to the PEP, which then denies or allows the action on the resource the original requester wanted (6). It is worth noting that the PEP and the PDP might be distributed across several servers [12]. The current version of XACML is 2.0.

3.3.2 Defining Policies

Access policies are defined using the XACML policy language. The root of all policies is a PolicySet which is a container that can hold other PolicySet elements and Policy elements, as well as references to policies found in remote locations. The Policy represent a single access control policy through a set of Rules, in addition to a Target element.

Since a PolicySet can contain multiple policies or rules, each of which may be found applicable to an access decision and evaluate to different access control decisions, XACML needs some way of evaluating these answers and reach one single decision. This is done with the use of a set of Combining Algorithms, which represent different ways of combining multiple decisions into one single decision. Combining Algorithms can be used on two levels, the Policy Combining Algorithms used on PolicySet elements, and the Rule Combining Algorithms used on Policy elements. There are in total six algorithms defined in the 2.0 standard [11], but with the possibility to build custom algorithms, highly complex policies can be created.

The Target in the Policy element is used to find the right policy to a given request. To do this it contains a set of conditions for the Subject, Resource and Action. Boolean functions are

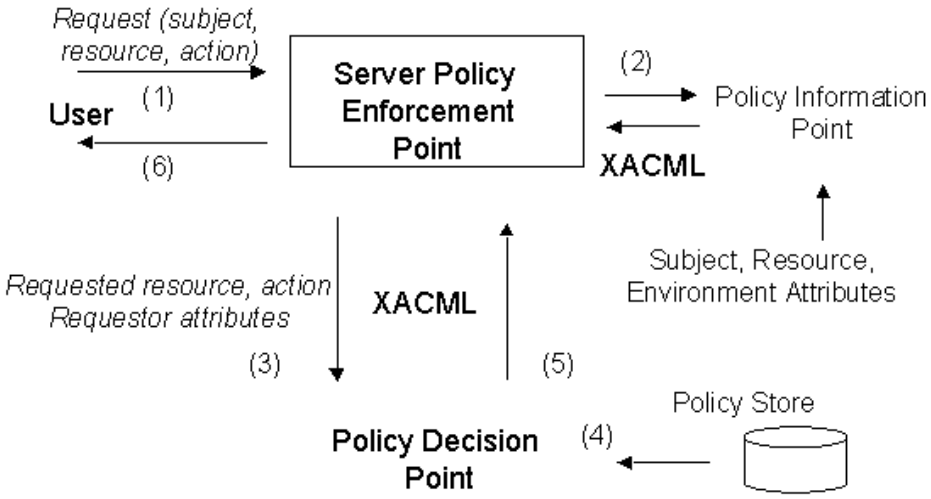

Figure 3.4: Data-flow diagram. Taken from [13]

then used to compare values from the request with the Target of existing PolicySets, Policies and Rules to find those who match and applies to that request. The Target is also used to index policies, making it efficient for lookup.

When the right policy has been found, its rules are evaluated. Each rule contains a Condition which is a boolean function. If the Condition is true, the Rule returns the Rule's Effect which in turn is evaluated by a Combining Algorithm if more than one Rule apply. The Rule's Effect is usually set to Permit or Deny. However, if the evaluation results in an error it returns Indeterminate, and a NotApplicable if the Condition does not apply to the request.

In the request and policies, Attributes are used to contain named values of known types. An Attribute may include an issuer identifier, issue date and time. A user name, a role, a specific document, and the time of day are all attribute values. It is the Attributes that are used when comparing a request from the PEP with a policy in the PDP.

The Policy has two mechanisms, the AttributeDesignator and the AttributeSelector, to resolve attributes from a request or some other source. The AttributeDesignator lets the Policy specify the attribute values by defining a name and type, and optionally an issuer as well, which is used by the PDP to look up values in the request. There are four kinds of AttributeDesignators, one for each of the type of attributes (Subject, Resource, Action and Environment) in the request. An AttributeSelector is used to define an XML Path Language (XPath) query where the attributes values should be looked for.

Because both the AttributeDesignator and the AttributeSelector can return multiple values, they are placed in a Bag. The values in the Bag then have to be compared to expected values to be able to make an access decision. This is done by the use of Functions. These Functions can be highly complex and nested, and it is also possible to write custom functions.

3.3.3 Defining Requests and Responses

In addition to defining how a policy should be built, XACML also has a standard way to express requests and responses. A request consists of Attributes in the following categories:

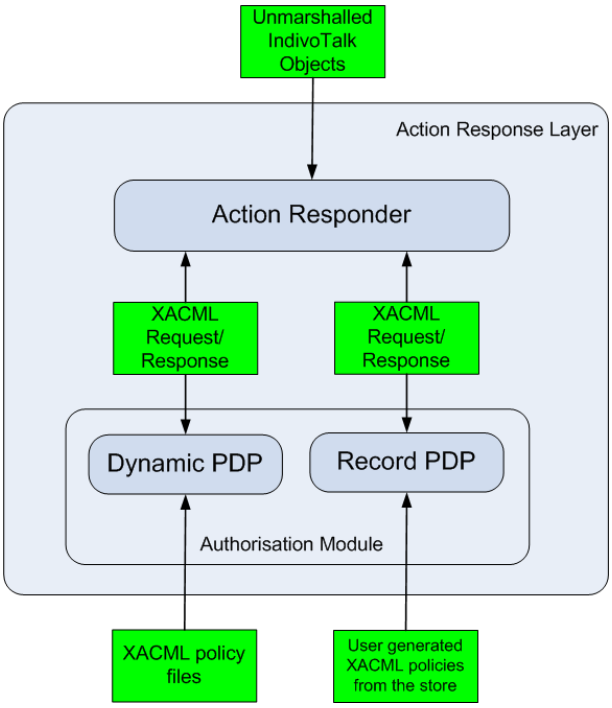

Figure 3.5: Action Response Layer

Subject, Resource, Action and Environment. The last category is optional. The Subject can consist of several attributes separated by different attribute ids, while the other categories can only have one attribute.

A response contains one or more Results, each consisting of attributes in the following categories: Decision, Status and Obligation. The last category is optional. Each Result represents the decision of an evaluation.

A XACML request, response and an example policy is shown in Appendix B.

3.4 Access Control in Indivo

This section explains more in detail how the access control in Indivo 3.0 Beta [14] is organised. Indivo is an extensive system, and we therefore do not go into detail in all parts, but will highlight the most important parts of the functionality regarding access control. Indivo utilises Sun's java based XACML implementation version 1.2 [15].

3.4.1 Action Response Layer

The heart of the access control in Indivo is situated in the Action Response Layer as shown in Figure 3.5. In this layer the different requests are interpreted and evaluated, and finally the responses are returned.

The Action Response Layer consists of a group of classes where each class is specialised to handle a specific type of action request, e.g. *DefaultAddDocumentActionResponder*, *DefaultRead-*

DocumentActionResponder and *DefaultUpdateDocumentActionResponder*. All these classes inherit key methods from a super class called *DefaultActionResponder*. Figure 3.6 shows a UML Sequence diagram of the process of authorisation of a request from when it enters the Action Response Layer to a response is prepared.

An incoming request is handled by the super method *processAction(...)*. This method unwraps the request and tries to authorise it using the *authorize(...)* method. The *authorize(...)* method is specially modified to the type of request, and is therefore defined in each of the ActionResponder classes. In addition to any specialised checks, the method also calls the super method *checkPermissionAndAudit* defined in DefaultActionResponder.

The *checkPermissionAndAudit* method generates a *RequestCtx* object, which is a Java representation of a XACML request containing information about the subject, resource and action among others things. This object is then sent to be evaluated by a set of system responder policies in the *DynamicPDP*, and a set of record specific policies in the *RecordPDP* in the two respective methods *isDeny(...)* and *isPermit(...)* as seen in Figure 3.6. A further explanation of responder policies and record policies can be found in Sections 3.4.2 and 3.4.3, respectively.

Results of the evaluation in the two PDPs are returned and combined with the results of the custom checks in the *authorize(...)* method, and together they return a deny or permit to the *processAction(...)* method. If the answer is permit, the request is processed in the *process(...)* method before a response is sent back to the requester.

Authorisation is thus done on two different levels in Indivo, the first following a set of special policies for the system, the second following policies based on the users own decisions. These are explained in more detail in the following two sections.

3.4.2 Global XACML Polices

An implementation of Indivo can have a set of global policies. This means that all requests have to be granted access based on these policies before any further evaluation. An example of such a policy can be to allow users logging in with the patient role to read and update his or her own record, but deny them the right to create new PHR records. The policies are stored in XML files located in the servlet's resource folder, and the policy module used by the DynamicPDP is created based on these files during initiation of the system. This means that any changes made to these policy files will not take affect before the servlet is restarted. An example of a policy file can be seen in Appendix B.1.

3.4.3 The Indivo Record and Record Specific XACML Policies

Each user in Indivo is assigned his or her own record. This record consists of several documents, e.g. discharge summaries, messages, annotations and lab tests. Each document consists of one or several document versions, and for each time a document is updated, a new document version is added. In that way the history of the document is kept, and a user can have a look at the version he or she wants. An illustration of the record structure is given in Figure 3.7.

The Indivo record consists of three system essential documents that are created automatically during the creation process of a new record. The three essential documents are:

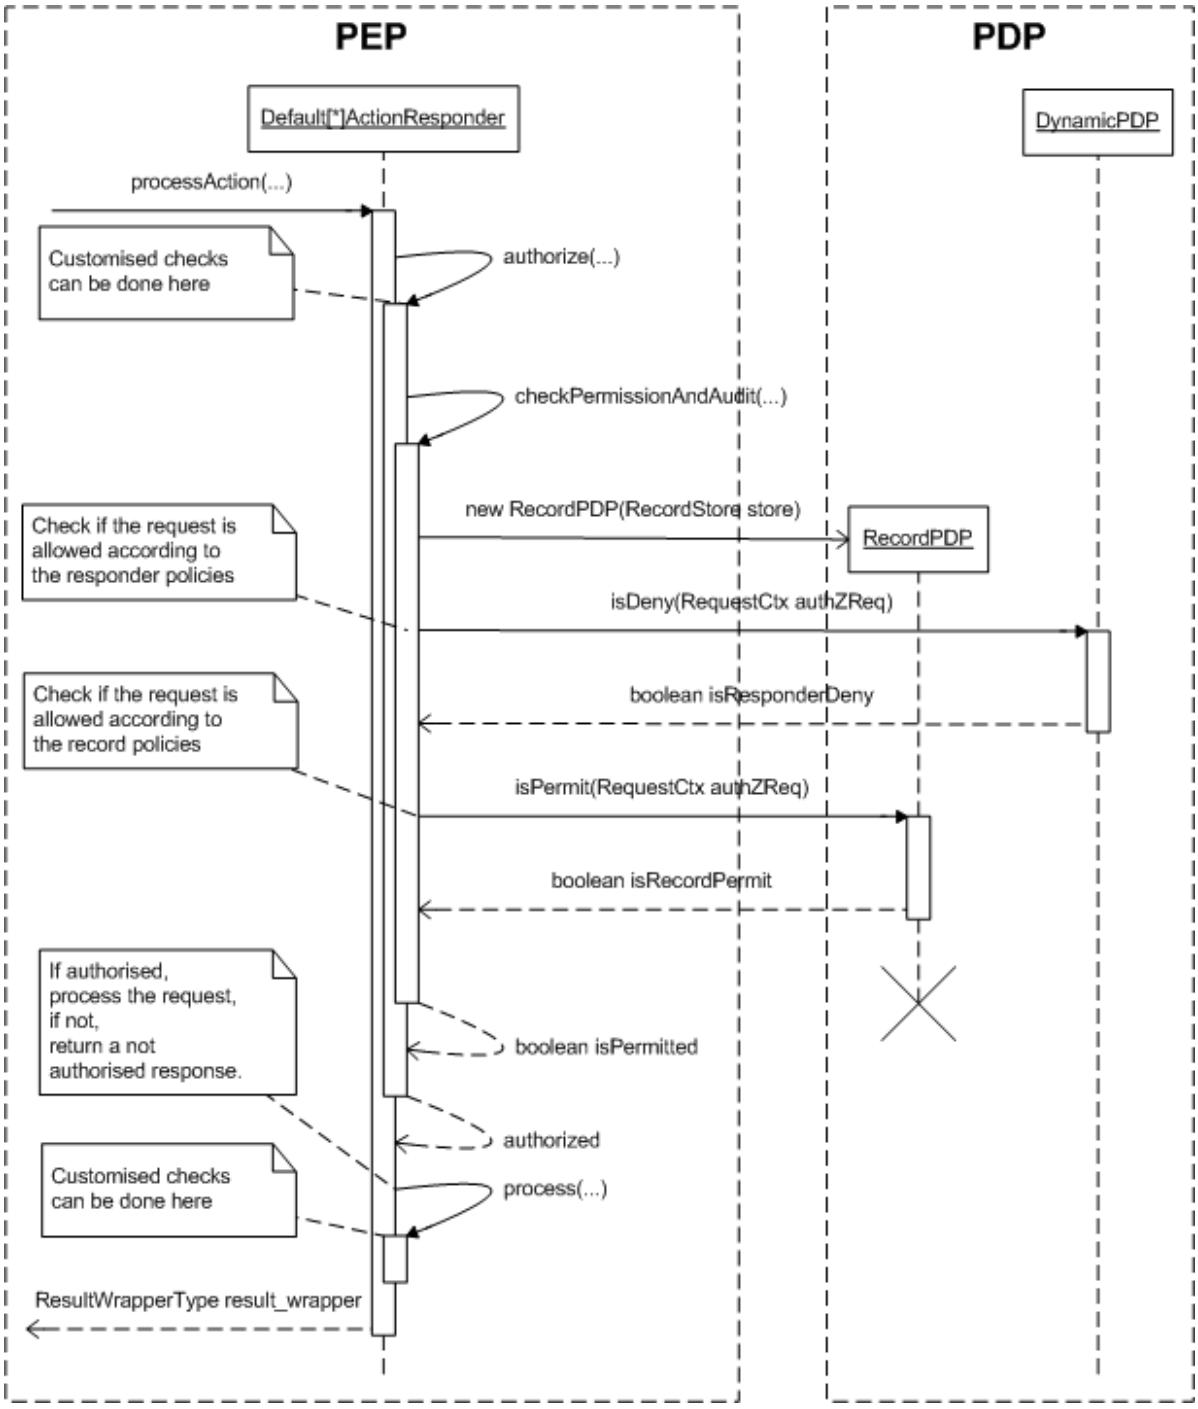

Figure 3.6: UML sequence diagram of Action Response Layer. [*] e.g. Add or Update

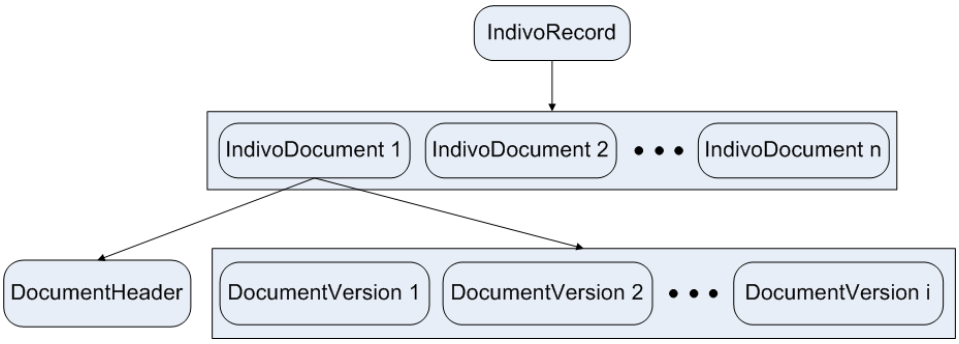

Figure 3.7: Indivo record structure

Credentials document. This document contains the password of the user, and is used during authorisation.

Actor attributes document. This document contains information about role and group membership of the user, as well as the user's default role and proxy representation. The roles in Indivo are associated with privileges by creating global policies as described in Section 3.4.2, while the group memberships do not have these privileges. They are, however, useful when creating record specific policies in that one can allow or forbid actions on part of the record to members of a particular group [10]. A proxy representation allows a user A to act as a proxy for another user B. This causes all the access policies on a record that apply to B to also apply to A.

Access policies document. This document contains the record specific XACML policies set by the user. It conforms to the XACML standard with a parent PolicySet, which includes a Target and one or several Policy and PolicySet elements. Each Policy defines a Rule for access based on a list of Subjects, a list of Resources, a list of Actions, and finally an Effect. The list of subjects is simply defining the users, roles or groups this rule apply to, the resources to the documents, the actions on the kind of action to be performed, and the effect is either deny or permit. These policies are updated in real-time when the user changes access control to any document in the record, and will take affect immediately.

CHAPTER 4

Previous Work

This chapter summarises our previous work done in the project "Access Control Model for Personal Health Record" [1], from now on referred to as the autumn project. The first section contains an overview of the project, and the next section deals with the interviews we performed. We also present our tentative access control model, before we round of the chapter with a summary of the project results.[1]

4.1 Overview

The goals for the autumn project was to start development of an access control model for the PHR based on principles according to the Norwegian health care system and its standards. Our concern was also to focus on the patient as a user of the PHR, and not only on the health care actors, who usually are the focus in health care informatics. To enhance this patient focus, we early in the project decided to carry out interviews with potential users of a PHR with the goals to get criticism and possibly verify the model.

4.2 Interviews

As mentioned we early on decided to use interviews with potential users of the PHR to achieve our goals. We were interested in achieving more knowledge of the needs from a patient point of view, as well as to, if possible, verify our access control model. In cooperation with Jorunn Bjerkan, a PHD student at NTNU, we carried out two pilot interviews which were used both as a foundation for our evaluation of the tentative access control model, and as preparations for further interviews. The methods used in the interviews and our preparations are presented in this section.

[1] Section 4.2 and 4.3 have been adapted from our autumn project "Access Control Model for Personal Health Record" [1]

4.2.1 Method

As preparations for these interviews we studied qualitative interviewing techniques. The book "Interviews. An Introduction to Qualitative Research Interviewing" by Steinar Kvale [16] was an important source and guideline for us when preparing for the interviews and when creating the interview guide. We also studied the possibilities to utilize participatory design, where the goal is to actively involve the end user, here the patient, in the design process.

Qualitative Interviews

Using qualitative interviews is a way of achieving new knowledge through conversations with other people. Kvale [16, p. 17] emphasises the meaning of the word interviews - *inter views* - an exchange of *view*points between persons that converse about a subject of shared interest. Qualitative interviews are typically only performed using a handful of participants. The interviews can be structured, semi-structured or more or less informal depending on what kind of knowledge that is wanted. It is usual to divide a qualitative interview study into seven stages [16, p. 46-49]:

1. Topicalisation. Formulating the purpose of the study and the main questions, and present a theoretical analysis of the subject that is to be investigated should be done in this stage. The questions *what*, *why* and *how* are central:

- *what*: come up with foreknowledge about the subject in question
- *why*: make clear the purpose of the study
- *how*: come up with knowledge about different interview and analysis techniques, and decide which technique to use when gathering the wanted knowledge.

2. Planning. In this stage the study is planned taking all seven stages into account. One should plan with an eye to collecting the wanted knowledge and towards any moral implications that might arise in the process.

3. The interview. One should carry out the interviews as planned in an interview guide, and try to collect as much relevant knowledge as possible while maintaining good inter-human interaction.

4. Transcription. In this stage the interview material is prepared for analysis. Usually this means transcribing from video and/or audio to text.

5. Analysing. On the basis of the purpose and the subject area of the study, and in accordance with the nature of the transcribed material, a method for analysis is chosen. The transcribed material is then analysed.

6. Verification. The findings from the interviews are in this stage examined regarding generalisation, reliability, and validity.

7. Reporting. In this stage the findings and the methods used are prepared in a form that satisfies scientific criteria and the study's moral aspects.

Participatory Design

Participatory design (PD) can be defined as "a set of theories, practices, and studies related to end-users as full participants in activities leading to software and hardware computer products and computer-based activities" [17]. Combining the knowledge of the end-user and the software developer, with practises neither in the end-user's domain nor in the developer's domain, results in a shared "in-between" region with shared attributes of both parties. This "in-between" region, or "third space", have shown to be valuable in combining knowledge into new insight [17]. Therefore, by involving the end-user actively in the development, the result can be better products and services. PD is very diverse and can be used in most areas in computer science. This diversity also leads to no single paradigm of study or approach to practise of PD.

An approach to PD is PICTIVE, as explained by Muller [18]. PICTIVE, or Plastic Interface for Collaborative Technology Initiatives through Video Exploration, is a technique that utilises low-tech objects in a kind of brainstorming session. These low-tech objects can be made out of plastic or other simple materials, such as notes, in different colours or shapes. By using such objects to represent system functionality, one ensures that all participants in the session have equal opportunities to contribute with their ideas [18]. Another important quality is that objects can, and should, be modified in real-time. Video recordings makes it easier to analyse the sessions.

Motivation

The qualitative interview is a good method to extract useful knowledge from the participants. They not only answer questions prepared by the expert, but also formulate their own understandings of the world they are part of through the dialogue with the interviewer [16, p. 25]. In that way the qualitative interview adds a new dimension of knowledge compared to e.g. a quantitative survey or no user interaction at all. This also makes the process of creating the model more challenging, since the knowledge has to be interpreted, analysed and verified in the right context before it can be used.

The reason for using participatory design as part of the qualitative interview, is to be able to describe our model in a non-technical fashion so that the participants can understand it. Traditional representations such as requirements specifications have shown to provide only limited support for participants who want to understand a new model. In order to fully understand it, the users have to get hands-on experience; to try it out and play with it [19]. Hence, by using participatory design we hope to create a rich and creative discussion and receive valuable feedback.

4.2.2 Preparations

One of the most important processes in a qualitative interview is the preparation phase. Without a proper grounding consisting of methods to use before, during and after the interview, the results might not satisfy what was expected.

The following subsections describe our preparations with the pilot interviews. They are loosely based on the seven steps described in Section 4.2.1.

Topicalisation

Our main motivation and goals for the interviews were to investigate whether:

- our access control model was usable and understandable for the user.
- our suggested roles and access matrix were reasonable, and either correct, reduce or expand these.

Further the pilot interviews would give us valuable experience with interviewing and the use of participatory design for future interviews.

The Interview Guide

Our interview guide outlined the subjects of interest:

1. Introduction
    - About us and our assignment
    - Goal of the interview
    - Plan for the interview
    - The PHR - what is it?
2. Roles
    - What persons/professions do you think can benefit by taking part in your PHR?
    - What kind of roles have you been in contact with?
    - What kind of roles do you know within the health service, and are these interesting in a PHR setting?
3. Content
    - What kind of content do you think it can be interesting to share in the PHR?
    - Have you ever looked in your record and what kind of documents was this?
    - What kind of functionality would you have in PHR
4. Access control model
    - Short presentation of our model
    - Discussion about the model, thoughts, reasonable, missing anything, too complicated?
5. Access matrix - cases giving roles access to content
6. General comments and questions from the participant

The guide was supposed to work as a guideline for the interview, outlining only the subjects and some key questions. Through this guide we hoped that we would get feedback on our goals and motivations. The interview itself was semi-structured, which meant that supplementary questions would have to be asked. These questions were, however, not prepared in advance. Before the interviews we chose to prioritise item 2 and 4 in the interview guide in case we had too little time for all. We decided to conduct two interviews with one participant at a time. Each interview would be parted in two: Jorunn Bjerkan would first conduct her part, and then we would conduct ours.

Using Participatory Design

Our participatory design workshop would consist of the participant, an interviewer and an interview assistant. It was supposed to be an integrated part of the interview. The following items would be used:

- a table
- regular size post-it notes of two different colours
- small stickers of several different colours

One half of the table would be reserved for post-it notes representing resources, and the other half for post-it notes representing roles. If the participant would name relevant resources or roles during the interview, the interview assistant would write these down on the post-it notes and place them on the table. Later during the access matrix item in the interview guide, the participant would give the roles access to the different resources on the table using the small stickers. Each role would then get a unique coloured sticker, and if it was given access to a resource, an equivalent sticker would be placed on that resource. An illustration of this is given in Figure 4.1.

Transcription

Performing the transcription can be said to be an interpretation process in itself. The transcriptions are not mere copies or representations of the interview, they are abstractions just like topographical maps are abstractions of the real ground [16, p. 104].

We decided to use video/audio recording equipment during our interviews, and transcribe the recorded material. In this way we would be able to write down what the participants were saying and interpret their body language. We chose a writing style that lay as close up to the participants verbal language as possible.

Each transcriber would read through the other transcriber's work while watching the video in order to check the quality of the transcription.

Method of Analysis

When choosing a method of analysis, we wanted one that retained as much of the "whole" of the data as possible while at the same time managing to capture the participants meanings

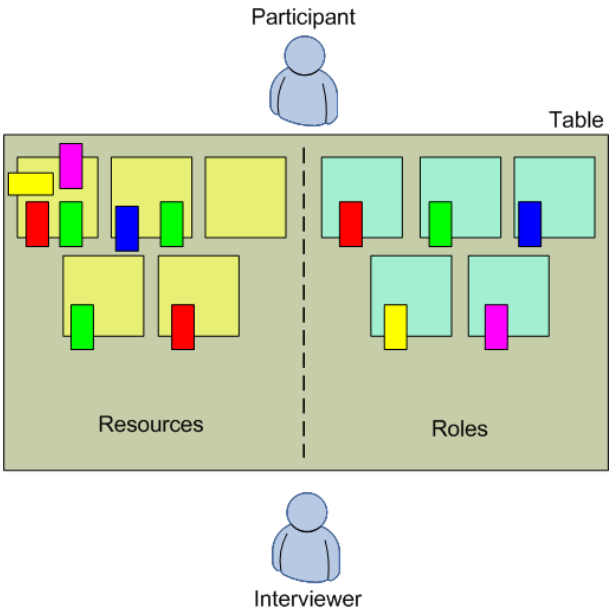

Figure 4.1: The participatory design workshop

and experiences. We ended up with a method inspired by "meaning condensation" [16, p. 121-129]. Our method consisted of five steps. The first step was to read the transcribed interview to get the big picture of the session. Then the interview was reread, but this time all the meaningful statements were labelled. In the third step the statements were studied and their meanings were expressed as simply and clearly as possible. Some meaningful statements[2] and their meanings are given in Table 4.1.

Statement	**Meaning**
I sometimes write stuff in my log that I don't want others to see. Or... I want some co-workers to be able to read it, but not everybody.	Wishes access control for individuals.
I think that giving access to everybody in my workplace is a horrible idea because... well... I don't trust everybody.	Does not want to give access to an institution.

Table 4.1: Text analysis, step 3

In the fourth step of the analysis we sat together analysing and discussing the statements and the meanings. This was done in the light of three of the main subjects from the interview: roles, record resources and the access control model. If relevant, the meaning would be placed under the correct main subject from the interview. E.g. in the above example the meanings would be placed as shown in Table 4.2.

The fifth and last step was to organise and tie all the meanings under each main subject together and write a continuous and meaningful analysis.

[2]The statements are fictional.

Main subject	Meanings
Access control model	Wishes access control for individuals.
	Does not want to give access to an institution.

Table 4.2: Text analysis, step 4

Verification

In a qualitative interview, time should be spent to verify the interviews with regards to reliability and validity. Due to time limitations and the fact that we were only performing two pilot interviews, this part of the interview process was omitted.

4.3 Access Control Model

This section describes our work and results of the developed tentative access control model.

4.3.1 Method

Work with the design of the tentative access control model was based on several techniques. The study of state-of-the-art literature about relevant subjects was one of them. This included important information such as the understanding of the basic principles of access control. Further, collected data from KITH standards, various EHR systems and their access control, and other projects with similarities to the PHR were an important source of information and inspiration.

The work was also greatly influenced by the concept of role-based access control (RBAC), which is a popular choice for access modelling. RBAC is based on assigning users to roles and giving different permissions to different roles, and Ferraiolo et al. have proposed a standard for such a solution [20].

Discussions and input from others were a great source for new ideas. The development of the tentative access control model was based on iterative work, and it was valuable to receive other people's opinions and criticism in the process of improving the model.

4.3.2 Preparations

As preparations for our work with the access control model we studied access control in general. Further we looked into access control in existing EHR's in Norway. This included Siemens' DocuLive ERP, DIPS ASA's DIPS and TietorEnator's Infomedix. We also examined the two projects The Kjernejournal and SamPro, which are both ongoing and have similarities to the PHR.

A part of our work was also to suggest roles and institutions that could be used in the access control model. In addition, we came up with a suggestive list of possible resources that the PHR should contain. For both these purposes we used standards developed by KITH as basis.

4.3.3 Results

This section describes the tentative access control model we developed during the autumn project.

Basic Keywords

The use of our model can be divided into three basic keywords:

- Subject
- Record resources
- Operations

The subject defines who should have access in an access control model. Our model is influenced by the principles of RBAC, which only use the role concept as a subject. Because we believe it is too restrictive and narrow to only use the role concept in a PHR we have expanded the model to also handle individual users and institutions in addition to roles. This makes access control possible on various levels of granularity, depending on the patients' wishes. Our suggested list of roles and institutions participating in the PHR can be found in the autumn project.

The concept of record resources deals with the concrete information that is being handled in the access control. For access and authorisation control to be valuable the PHR needs to be divide into useful parts. This enables the authorisation mechanism to give access to specific parts of the record instead of the whole record or only single documents. Since there does not exist any standard for PHR content and record structure we came up with a suggestion for this. The PHR record structure we suggested was based on KITH's EHR standard, which divides content in groups of cases, documents or fragments. We chose to omit the use of fragments in our access control model because we believe it would complicate the management of access control. The suggested list of PHR content can be found in the autumn project.

Operations defines the actions a user can perform on a resource depending on the access control settings. When setting access control the user can give another user, role or institution one of the following operations:

- No access
- Read access
- Read and write access

Access Control Model Formalism

The access control model is based on the terms subjects, resources and operations defined in the previous section. This section explains the access control model developed in the project. Our model is based on work as described in Section 4.3.1.

The model is depicted in Figure 4.2, and shows the data elements and relations that defines the model. It is important to understand that in the model, access is given to resources in one specific record, and not to resources in general. This is shown in the figure as two separate views of the model, the global view for the whole system in Figure 4.2(a), and the record specific view in Figure 4.2(b). The data elements consist of USERS, ROLES, INSTITUTIONS, INSTITUTION-ROLE ASSIGNMENTS, RESOURCES, OPERATIONS, and PERMISSIONS. All described in Section 4.3.3, and where the three first data elements correspond to the subject.

The model also illustrates the relations depicted as arrows between the data elements. PERMISSIONS are the result of the relation between RESOURCES and OPERATIONS. Access is then given based on rules formed by permissions given to a subject. The figure shows the many-to-many relation between resources and operations which form the permission that denotes an approval or denial to perform an operation on one or more resources.

The *global user assignment* (GUA) and the *record user assignment* (RUA) relations both show that a user is assigned one or more roles, and a role can be assigned to one or more users. While the GUA applies to the whole system, seen in Figure 4.2(a), the RUA only applies to one specific record, seen in Figure 4.2(b). In this way roles that are specific for each patient, e.g. primary physician, can be mapped to a user in each record using RUA, while a user can be mapped to the role physician throughout the whole system using GUA.

The *role assignment* (RA) relation shows a many-to-many relation between roles and institutions, meaning that a role can be present in many institutions, and that one institution can have many different roles assigned. This relation forms the new data element INSTITUTION-ROLE ASSIGNMENTS which is used when assigning access for roles and institutions.

In Figure 4.2 we can see the permissions assignments relations between the subject, top, and the permissions, bottom. The *user permission assignment* (UPA) is used when the subject is a user; access is then given to a specific user. To be able to assign permissions to roles in general, and to roles affiliated a specific institution we have introduced the *institution role permission assignment* (IRPA). IRPA is a combination of a role permission assignment (RPA) and an *institution permission assignment* (IPA) denoted by dotted lines. RPA gives a role at any institution permission, while IPA gives an institution permission. Because we do not want to give everybody at the institution the same access, but rather assign different permissions to the roles affiliated to the institution, the relation between roles and institutions have to be used. The model solves this through the use of the IRPA relations which utilise INSTITUTION-ROLE ASSIGNMENTS.

We can therefore denote the following rules for defining access to a record:

- `UPA: (P, U) where P is a permission and U a user.`
- `IRPA: (P, (R, I)) where P is a permission, R a role and I an institution.`

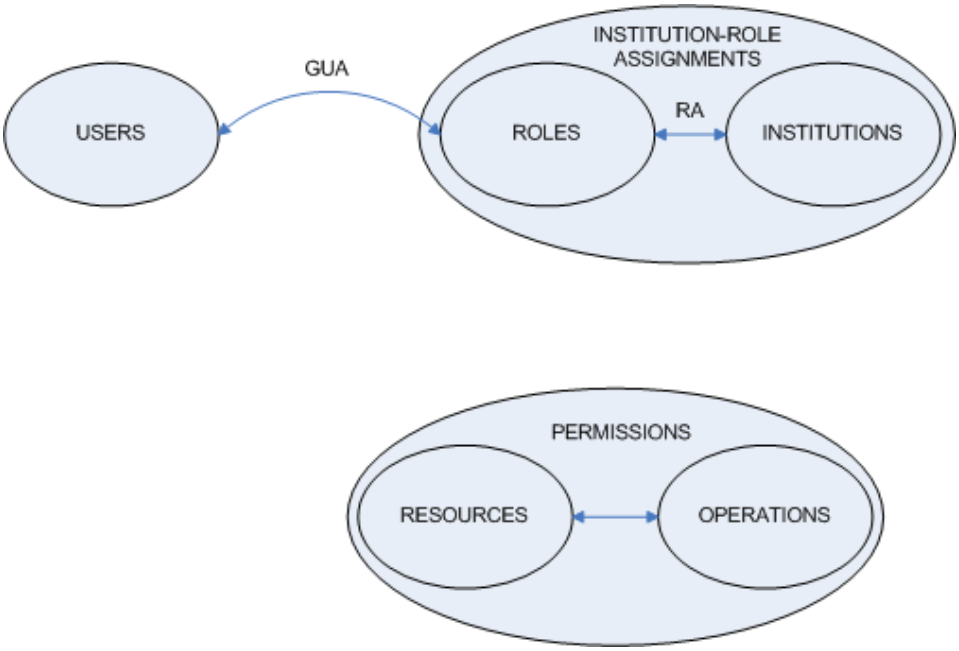

(a) Global specific view

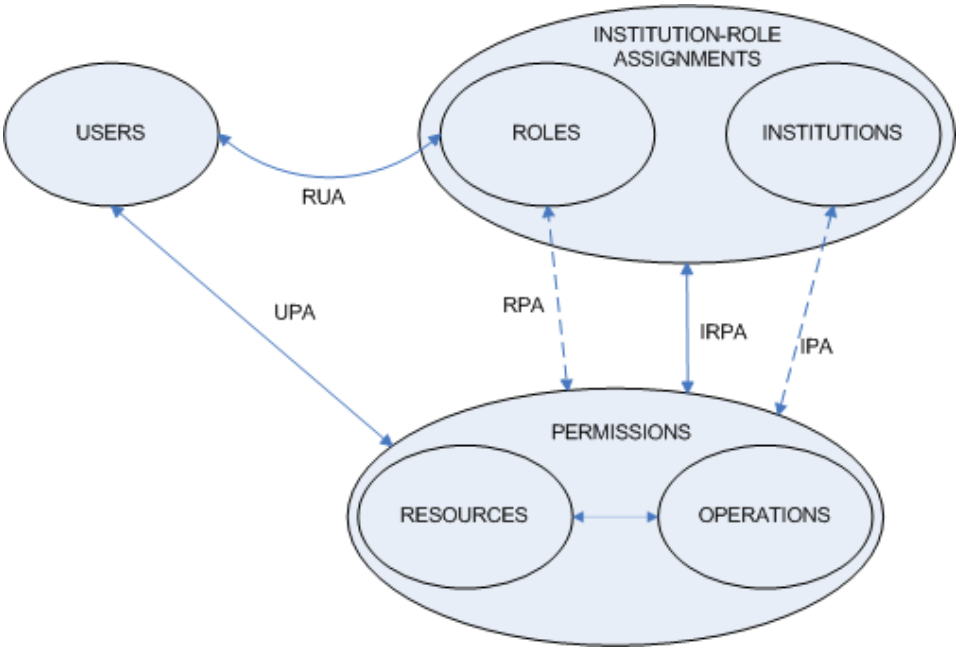

(b) Record specific view

Figure 4.2: General access control model: (a) Global specific view and (b) Record specific view

- `By setting R as wildcard _ we get (P, (_, I)) which denote that P is assigned to all roles affiliated to institution I`
- `By setting I as wildcard _ we get (P, (R, _)) which denote that P is assigned to all roles disregarding affiliation to an institution`

This ability to give permissions to both users and institutions in addition to roles, provides great flexibility and granularity. It will enable the owner to adjust access from strict, giving permissions to named users, to a more open and general access using roles and institutions.

By using the hierarchy of resources we gain even more flexibility and granularity in our model. Figure 4.3 shows this hierarchy, where the term resource is a generalisation of a case or a document. Operations are related to these different levels, which gives the ability to define access to whole cases or only documents.

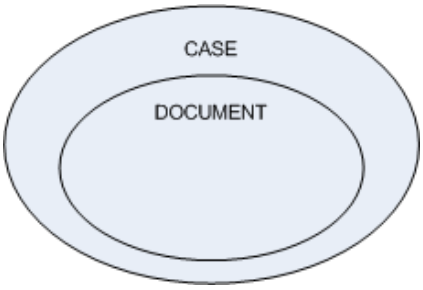

Figure 4.3: Resources

The following is a formal summary of our model:

- *USERS, ROLES, INSTITUTIONS, RESOURCES,* and *OPERATIONS*, basic data elements.
- $GUA \subseteq USERS\ x\ ROLES$, a many-to-many mapping user-to-role assignment relation for the whole system.
- $RUA \subseteq USERS\ x\ ROLES$, a many-to-many mapping user-to-role assignment relation for a specific record.
- $assigned\_users\_global(r) = \{u{=}USERS \mid (u,r) \in GUA\}$, global system mapping of role r onto a set of users.
- $assigned\_users\_record(r) = \{u{=}USERS \mid (u,r) \in RUA\}$, record mapping of role r onto a set of users.
- $assigned\_roles(i) = \{r{=}ROLES \mid (r,i) \in RA\}$, mapping of institution i onto a set of roles.
- $PERMISSIONS = 2^{(OPERATIONS\ x\ RESOURCES)}$, the set of permissions.
- $INSTITUTION\text{-}ROLE\ ASSIGNMENTS = 2^{(INSTITUTIONS\ x\ ROLES)}$, the set of institution-role assignments.
- $UPA \subseteq PERMISSIONS\ x\ USERS$, a many-to-many mapping permission-to-user assignment relation.

- *IRPA* $\subseteq$ *PERMISSIONS x (INSTITUTIONS x ROLES)*, a many-to-many mapping permission-to-institution-role assignment relation.
- *assigned_permissions(u)* = $\{p \in PERMISSIONS \mid (p,u) \in UPA\}$, the mapping of user u onto a set of permissions.
- *assigned_permissions(i,r)* = $\{p \in PERMISSIONS \mid (p,(i,r)) \in IRPA\}$, the mapping of institution i and/or role r onto a set of permissions.
- *Ob(p:PERMISSIONS)* $\rightarrow \{op \subseteq OPERATIONS\}$, the permission-to-operation mapping giving the set of operations associated with permission p.
- *Ob(p:PERMISSIONS)* $\rightarrow \{re \subseteq RESOURCES\}$, the permission-to-resource mapping giving the set of resources associated with permission p.

4.4 Autumn Project Results

This section presents a summary of the results from the interviews, and a summary of the work with the access control model.

4.4.1 The Interviews

The two pilot interviews gave us valuable experience. We found out that our main structure of the interview guide was satisfactory, but that some of the points needed changes. The participants did not have the expected knowledge about the Norwegian health care system, and they also tended to focus their answers around their own case history. Hence, some of our questions were not good enough.

We also experienced time issues. Our interview guide was too comprehensive compared to the time available, and we did not manage to finish all the points. This especially affected the PD part of the interview, as this was our last point in the guide. In addition, the PD was not prepared well enough.

The interviews also gave us feedback on the access control model, both negative and positive. In general, the participants approved our work, as the access control model supported the participants' wishes.

The results showed us that more work had to be done to the interview guide and to the use of PD.

4.4.2 The Access Control Model

The development of the access control model was an iterative process throughout the autumn project, as we got new ideas and became aware of problems. As previously mentioned, we received positive feedback on the model through the two pilot interviews, but the model should be further verified with new and improved interviews.

Among the issues that emerged was the lack of a Bluelight function, and the lack of role and institution hierarchies. We also realised that we needed to work more on our role list, the access matrix and the PHR resource content.

We believed, however, that the foundations for a good access control model for the PHR had been laid.

CHAPTER 5

Interview

In our autumn project "Access Control Model for Personal Health Record" [1], we conducted two pilot interviews in order to get feedback on our design of the access control model. Although these interviews provided valuable results, it was always our intention to arrange more interviews in order to get a wider basis to draw conclusions from. This chapter describes how we have prepared for the new interview round. For an introduction to qualitative interviews, the reader is referred to Section 4.2.1.

5.1 Interviewers and Interviewees

Our pilot interviews from the autumn project were conducted together with PhD student Jorunn Bjerkan. Through her project leader role in the Sampro project [8], she had access to both patients, IP coordinators and health care actors. This co-operation was continued. In addition, another master's degree student joined us, which made the interview team consist of 4 people in total.

One weakness with our previous work was that we only studied the access control model from the patient's point of view. For our access control model to succeed it should be accepted by all users, and hence health care actors should be included in our study. Sitting at "the other side of the table" as opposed to the patients, health care actors probably has valuable feedback regarding subjects that patients do not think about. We have decided to address this issue by including IP coordinators, health care actors and patient ombudsmen in our interviews. A patient ombudsman provides legal protection for the patient by helping him or her to obtain health care services, and by assisting the patient in complaints and lawsuits against the specialist health care services. He also contributes to the general improvement of quality in the specialist health care services, mainly as a result of the contact with the patients. The health care actors and some of the IP coordinators we will interview are employed within

psychiatric health care, special pedagogics, physiatrics or the social sector. Two of the patient interviews will be performed with the patient guardian instead of the patient.

All the interviewees except the patient ombudsmen were involved in one or several IPs. Some of them were also working on the same IP. An overview of the participants is given in Table 5.1. All participants except the patient ombudsmen have been made anonymous.

IP	Role
1	Patient IP coordinator Health care actor
2	Patient guardian IP coordinator Health care actor
3	Patient guardian IP coordinator Health care actor
4	IP coordinator Health care actor
	Nils Hybertsen, patient ombudsman Sør-Trøndelag Kjell J. Vang, patient ombudsman Nord-Trøndelag

Table 5.1: Interviewees

The interviews will to be carried out with one interviewee at a time and three or less interviewers using a digital video recorder. The restriction of interviewers are introduced to avoid an uncomfortable interview setting for the interviewee.

5.2 Interview Goals

The goals with the interviews have not changed much compared to the goals with our pilot interviews. However, we have decided to leave further work on the access matrix, and focus exclusively on the model and the most important PHR roles. Creating an access matrix with satisfactory quality will be too time consuming in this thesis, and it will also require separate interview workshops with patients and health care actors. For the same reason our work with concrete institutions in the interviews has been left out. Our goals are shown in Table 5.2.

Goal	Description
1	Investigate whether our access control model is usable and understandable in a PHR context
2	Identify the most important roles that should be present in a PHR

Table 5.2: Interview goals

5.3 The Interview Guide

Since we have decided to stop working on the access matrix, all of the points regarding content in a PHR and the access matrix have been removed. We hope that by doing this we will get answers of higher quality since there is more time for each remaining point. The revised interview guide is shown in Table 5.3.

In the introduction phase we have decided to focus on a more thorough explanation of the PHR by using sketches and images to illustrate the concept. We believe that the combination of verbal explanation and visual cues and illustrations will lead to good response from the participants. By further giving the interviewee as much knowledge about the PHR as possible, we hope to get better answers later in the interview.

In the role phase we have changed some of the content. Instead of asking direct questions like "what roles would you like to have in a PHR", we want to use a more indirect approach. By doing this we hope to get a fruitful discussion going and pick up important roles along the way.

The main change in the access control model phase is the way we use participatory design (PD). This is further explained in Section 5.4. We have also decided to give a more thorough explanation of the model for the same reason as for the explanation of the PHR in the introduction phase. Our PD cases are patient centred and will only be fully performed with interviewees that are patients. However, the PD material will also be presented to the other interviewees, and small examples will be given.

The interview guide is still semi-structured, where a shared understanding of the situation is a goal. This is done by going into a dialogue with the interviewee rather than maintaining a hierarchical structure, where the interviewer is the active part and the interviewee is passive. The interviewer starts the interview with an open mind and communicates in an co-operative way. While the interview is in progress, the interviewer interprets, encourages the interviewee, and asks new and improvised questions based on the received answers [21, p. 248-250].

As with the pilot interviews, Jorunn Bjerkan and the master's degree student will conduct their part first, and then we will conduct ours. Estimated total time of each interview is 1 hour.

5.4 Using Participatory Design

In the autumn project we experienced some problems related to the PD part of the interviews. We did not have enough time and the whole workshop should have been better planned. To address these issues we have improved our PD part significantly. In order to test our model we have introduced laminated cards representing the different units you can give access to and the units that make up a PHR. By using post-it notes of different colours you can link users and record content, and hence create access policies. In order to keep the PD workshop simple and less time consuming, we have decided to introduce an example with a fictional patient and some cases where he has to use the access control model. The interviewee is expected to act on behalf of this patient and perform his tasks. Only relevant PHR elements regarding the cases have been created. Table 5.4 shows the different elements that make up the PD workshop, and the different cases are as follows:

Interview Phase	**Asked to**	**Interview Points**
1. Introduction	A[a],B[b],C[c]	1. Presentation of the interviewers 2. Goals with the interview 3. How much knowledge about computers do you have, and what do you use them for? 4. Explanation of the PHR
2. Roles	A	1. PHR in your situation. Who do you share information with? 2. Other roles for other situations? Are you familiar with other chronic diseases? How would such patients use a PHR? Roles?
	B	3. PHR in your situation: In what way could you have used it?
	C	4. How do you think patients could have used a PHR?
	B,C	5. Which other participants/roles within the health care service would you have liked to have a closer co-operation with? In what way? 6. Bluelight functionality, who should have access?
	A,B,C	7. Do you know any situations in today's health care that are unnecessary complicated and that can become simplified by using a PHR? Which actors are involved?
3. Access Control Model	A,B,C	1. Explanation 2. General questions, thoughts, is there something missing, too complicated, is it possible?
	B,C	3. Quick examples using the PD cases 4. Sensible to split resources in document/case?
	A	5. Perform PD cases 6. Other points of view?
4. Wind-up	A,B,C	1. Questions or general comments?

Table 5.3: Interview guide

[a] Patient
[b] Coordinator/Health care actor
[c] Patient ombudsman

1. **Give a user access to a case**
   Mr. Krank is asked by his primary physician, Dr. Frisk, to give him access to the diabetes case in his PHR.

2. **Give an institution access to a case**
   Mr. Krank has to take insulin regularly, and he gets his medicine by showing his prescription at the pharmacy. To make this process easier the primary physician of Mr. Krank has started to publish electronic prescriptions in the PHR. If Mr. Krank wants his medicine, he now has to allow the pharmacy access to his electronic prescriptions.

3. **Give a specific user access to a specific document**
   Mr. Krank has to go to a diabetes specialist, Dr. Hansen. The primary physician of Mr. Krank writes a referral to this specialist and adds it to the PHR. Mr. Krank should now give Dr. Hansen access to this document.

4. **Give access to a role while denying access to a user with the same role**

   (a) Dr. Frisk and Mr. Krank agrees on an arrangement that if Dr. Frisk is sick, then Mr. Krank may contact one of the other physicians at the medical office. Mr. Krank wants to give all the physicians at the medical office access to the diabetes case.

   (b) After closer consideration Mr. Krank finds out that one of the physicians at the medical office is Dr. Sleip. Mr. Krank does not like this person, and does not want to give him access to the diabetes case.

Category	**Elements**
Institutions	Hospital, medical office, pharmacy
Roles	Primary physician, physician, coordinator
Users	Kåre Krank Dr. Frisk, (primary) physician, medical office Dr. Sleip, physician, medical office Dr. Hansen, physician, medical office Dr. Gundersen, physician, hospital Dr. Andersen, physician, hospital
Cases	Contact information, clinical history, diabetes, test results, individual plan, prescriptions
Documents	blood test, note, vaccine, physician referral

Table 5.4: PD elements

Figure 5.1 illustrates a solution to the last case, and also shows how the cards are designed. The yellow stickers indicate that all the physicians (Dr. Frisk, Dr. Sleip and Dr. Hansen) at the medical office have access to the diabetes case, which at the moment include a blood test document and a note. There is no need to put stickers on the 3 physicians or on the documents, since a sticker on a lower card (in this case the role physician and case diabetes) implies that the same access policy applies to all cards above. The red sticker on Dr. Sleip and on the NO spot in the diabetes case is an exception to the previous policy and means that Dr. Sleip does not have access to the diabetes case. Since there are no cards on top of Dr. Sleip, this policy only applies to him.

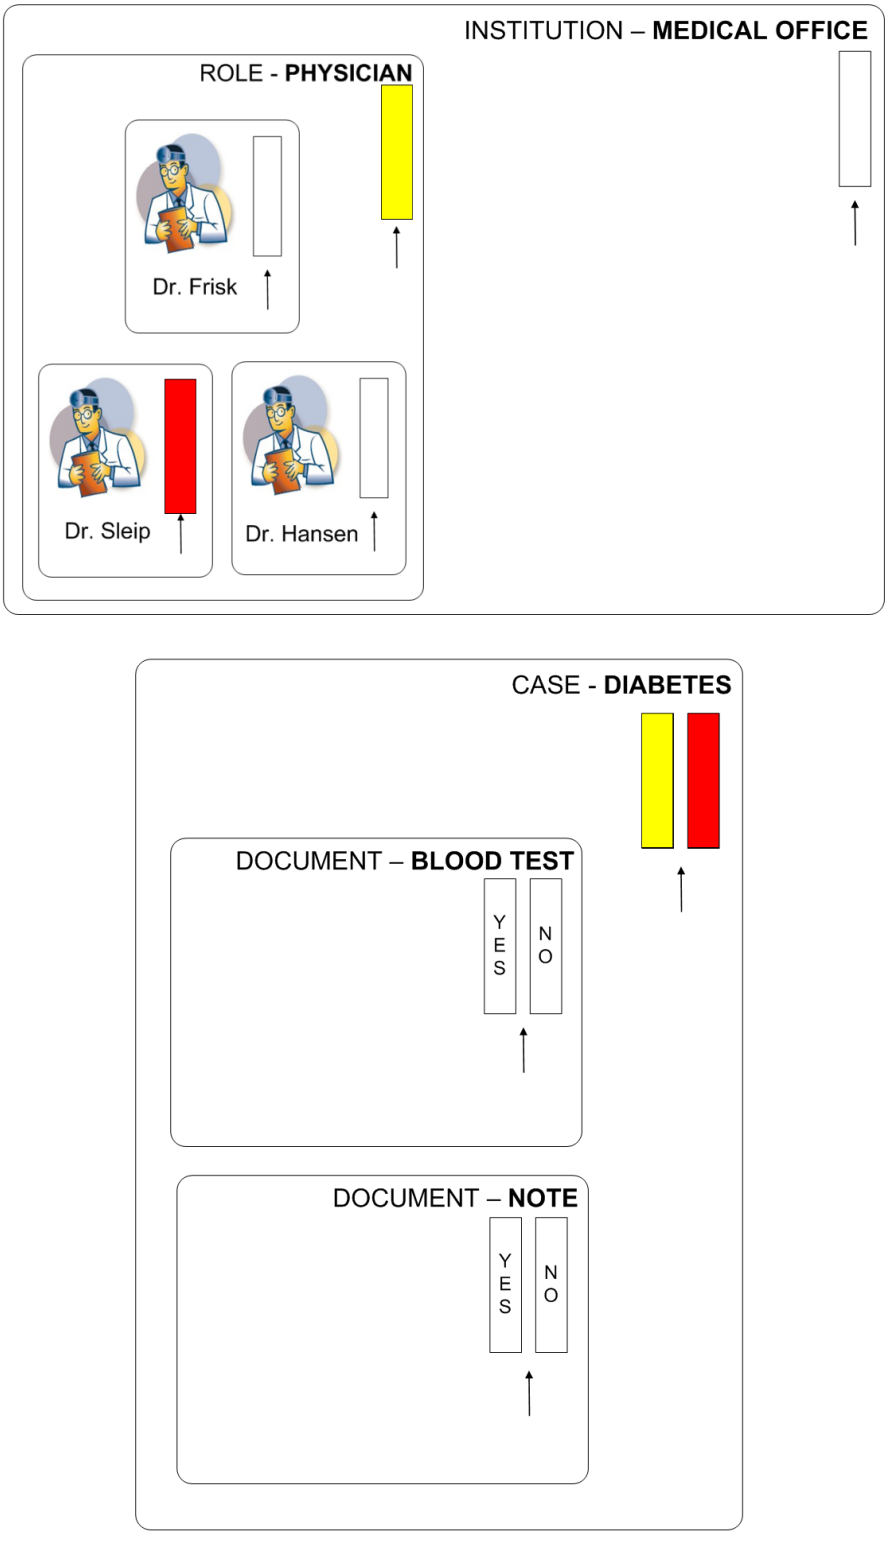

Figure 5.1: Example using PD

5.5 Analysing the Interviews

We have decided to transcribe the recorded interview material using the same method as with the pilot interviews described in Section 4.2.2. This technique has proved to ensure the necessary quality, and we feel there is no need to change it.

When it comes to the analysis of the transcriptions, we have made some minor changes. Since our collection of transcribed data will be much larger, we will use a software analysis tool called Nvivo [22]. Using Nvivo we will import the interview texts and be able to tag interesting statements and group text with similar meaning together, so that it will be easier to retrieve and compare. Our method is based on several known methods [21, p. 288-290][23], and it is still very similar to the one described in Section 4.2.2. It now consists of the following steps:

1. Read the transcribed interview to grasp its meaning as a whole. Try to be as open as possible.
2. Reread the interview and label the meaning units. A meaning unit is a part of the text that contains just one meaning, and can be everything from just a word to several sentences.
3. Study the meaning units and express their meaning as precise and clear as possible. An example showing a meaning unit and a meaning is given in Table 5.5.

These first three steps are made with a paper copy of the interview. The following steps are made using Nvivo.

4. Sit together and discuss which meaning units should be tagged. Locate the chosen meaning units in the imported interviews in Nvivo and tag them with their essential meaning. If there already exists a similar meaning, use that one instead. That way similar meaning units are grouped together.
5. Collect all meanings and group them under the correct sub-theme. Examples of sub-themes can be ”the access control model” or ”today’s health care system”.
6. Write a continuous and meaningful analysis based on the collected and analysed material.

Meaning Unit	**Meaning**
I sometimes write stuff in my log that I don’t want others to see. Or... I want some co-workers to be able to read it, but not everybody.	Wishes access control for individuals.
I think that giving access to everybody in my workplace is a horrible idea because... well... I don’t trust everybody.	Does not want to give access to an institution.

Table 5.5: Analysis of text, step 3. Taken from [1]

The analysis can be verified in two ways: by using several interpreters, and by accounting for how the analysis is performed [16, p. 136-137]. Regarding the first point our transcribed interview material will be controlled by a second interviewer against the recorded video. Verification also happens in step 4 and 5 in our method, since we perform this step together. Regarding the second point, we have in this section given a description of the method we use, and provided a small example. We will also provide some central quotations in the analysis chapter.

CHAPTER 6

Analysis

This chapter contains the results of our interview analysis. The first section gives a brief description of how we carried out the interviews. The second section deals with interesting topics that not necessarily were in our interview guide, but still became a subject of conversation. Existing problems in the health care system, and thoughts around the PHR in general are among them. The two sections after that deal with roles and the access control model respectively. The chapter is concluded with the findings from the patient ombudsmen interviews.

6.1 Accomplishment

All the interviews were completed over a period of two months. They were set up at the interviewees workplace, treatment centre or home residence, and recorded using a video camera. One of the interviews were performed using telephone. The average interview time was approximately 1 hour and 15 minutes, where our part took about 40 minutes. The PD workshop was fully performed with only one interviewee, due to time limitations in the other interviews. When the interviews were finished, the video recordings were digitalised and transcribed.

Both the video recordings and the transcriptions are confidential information. They are therefore not included in this report. The analysis is, however, not considered to be confidential. The quotations presented in this chapter have been made anonymous and are translated from Norwegian.

6.2 General Results

One of the first questions we asked the interviewees was how they considered their skills with computers and what they used it for. A great majority reported that their computer skills

were on a level where they could perform necessary work tasks. There was little interest in learning new skills that they would not benefit from in their daily life:

> **Interviewee:** *My view on this issue is that I learn what I need to know, because I won't take the trouble to sit next to the computer and learn stuff that I know I will never need. So when I wanted to learn how to use Sampro, then I learned it.*

In addition to using Sampro, work tasks performed typically included word processing, sending mail and using various work specific software tools. Most of the interviewees reported that they at home used internet for a few basic tasks like accessing their internet bank and writing e-mails. One interviewee had, however, successfully bought art on the internet. Another reported having little knowledge about how to use a computer.

During the interviews several interviewees started talking about weaknesses with the health care system today. Most of these issues were related to poor communication. The patients reported that it was cumbersome to make appointments with health care actors, and that it was difficult to get the information they wanted. The health care actors realised that communication among themselves could and should be more efficient than it is today, both regarding the patients, and their own efficiency. The following quotation is from a patient about poor communication between health care actors. Her medication received from home care services had not been updated after a hospitalisation:

> **Interviewee:** *Yes, because I have experienced that when I have received my pill dispenser, it has not... the medicine has not been changed.*
> **Interviewer:** *So the medicating has not been correct?*
> **Interviewee:** *shakes the head* *No.*

A possible problem reported by an interviewee, was that many patients probably did not feel that they had control over their own treatment. This was confirmed by another patient with a chronic disease.

Another main subject that was discussed during the interviews, was positive and negative aspects regarding the PHR. Some of the health care actors pointed out that the PHR should be efficient and user friendly. Health care actors would not want to use a system that slowed down their work routines. In addition, many of the patients who would need a PHR would also be the ones who would find it most difficult to use it, because of psychological, physiological and/or technical issues. This could in some cases be resolved by letting a trusted guardian take control over the patients PHR.

Especially the patients pointed out that a PHR would probably simplify communication between health care actors, and between health care actors and patients. This could, according to some of the health care actors, also lead to two unwanted situations. First of all, by allowing the patient to contribute to the PHR, one would risk that the PHR would fill up with a lot of meaningless patient created text that would make it more difficult to find relevant information. On the other hand, if the patient provided text would be relevant, it could contribute to an improved treatment. Secondly, there was a risk that one would most of the time communicate using the PHR, loosing a lot of the personal contact with the health service. The following health care actor describes such a problem:

> **Interviewee:** *But then I also wonder how fast that communication is. If a physician reviews a patient for.. lets say cancer, which is a serious diagnosis. Can he contact the patient and have a good dialogue before they get the review there* (in the PHR)*? I am thinking a little bit in proportion to reactions and such.*

The health care actor points out that one could experience situations where serious patient related information would be published in the PHR, and hence be made available to the patient before one could arrange a face-to-face meeting with explaining and support.

We came across an interesting find when talking to a patient about Sampro, which has much in common with the PHR. The patient reported that the health as well as the feeling of dignity and control over own life had increased after having started using Sampro as a IP tool:

> **Interviewee:** (about Sampro) *I think it gives us a more dignified life. Because it is like I have said, I don't want a paper based plan. It's no using forcing that on me.*
> (...)
> **Interviewee:** *Yes, but if they think about it, I haven't been admitted to a psychiatric ward in 2006. Normally I would have had 7-8 stays.*
> **Interviewer:** *7-8 stays in 1 year.*
> **Interviewee:** *Yes. That's how much money they have saved on one person, when thought about like that. And I reckon that there are lots of others.... who could have managed it* (Sampro) *and gotten a more dignified life. By just having that computer.*

6.3 Roles

One of the two main goals with the interviews was to identify the most important roles that should be present in the PHR. Table 6.1 shows the answers from the interviewees. The right column denotes in how many interviews the role were mentioned.

6.4 Access Control Model

In the interviews an explanation of our access control model was given, which led to discussions around the solution.

The overall reaction from the interviewees was positive, and three of them also stated that they did not see the access control model as complicated. However, there were also raised negative remarks towards the complexity of the access control model. These were made on the grounds that the model could be difficult for some patient groups, depending on age and illness. The following quotation shows this:

> **Interviewee:** *For some, now I'm thinking in general, for some patients it will surely be that* (the access control model is too complicated to control) *, but it depends on what the diagnosis is, what the issue is.*

Role	Mentioned in interviews
Primary physician	9
Physiotherapist	8
Ergonomist	7
Physician	6
Social worker	4
Psychologist	3
Parents	3
Community nurse, Special educator, Hospital physician, Dentist, Nurse, Psychiatric nurse	2
Spouse, Plan coordinator, Local case handler, Psychiatric community nurse, Public health nurse, Activator, Specialist physician, Pediatrician, Rheumatologist, Orthopedist	1

Table 6.1: Roles

When questioned if any help would be needed to manage the access control most of the interviewees clearly saw the advantage of this. This would also depend on the patient's own situation, knowledge and health status, some needing more help than others. The interviewees saw the need for a person with more knowledge that could be consulted, not only when administering access control, but also for the system in general:

> **Interviewee:** *Yes absolutely, I think people need a little guidance, a little guidance in the system, I will absolutely believe that.*

It was pointed out by interviewees that such help at least should be given in the start phase. When asked who should have the responsibility to offer the guidance, the answers were quite spread. Most pointed out that it should be a person that the patient had confidence in and had regular contact with. This person should also be well trained in the system, and have a good knowledge of the health care system. A person within NAV, the primary physician, or the coordinator in an individual plan were mentioned as possible roles for this task.

The interviewees had different opinions when we talked about giving access to roles and institutions in addition to specific users. Some felt that they could in some cases set access to roles and institutions, while others would prefer to set access only to individual users. One interviewee also expressed that he/she would want to have had personal contact with the user before giving access. The general opinion, however, was that the decision on who should have access, depended on the patient's situation.

Many of the interviewees expressed the need and importance to limit access to some of the information in their record. To be able to deny specific persons to access their record were in the interviewees interest, along with the ability to not share obsolete or sensitive information that was not relevant. A problem raised during some of the interviews was the danger of withholding important information from health care personnel. One of the interviewees also

raised the question of who should have the responsibility if a patient received wrong medical treatment, because the health care personnel did not have access to important information. Another interviewee expressed the danger of getting to restrictive, and that we, as patients, should trust the professional health care personnel:

> **Interviewee:** *(...) but I'm a bit afraid that we can become too restrictive as well, and that this can lead to us working against ourselves, (...)*

Another issue that was pointed out was the users ability for consents. Some might not understand the consequences and implications of their actions, even though they were healthy and sane.

In some of the interviews with health care personnel we discussed the use of the Bluelight function. The general opinion was that this was something that could be useful. The interviewees meant that such a function could be used by members of the emergency unit. However, one interviewee emphasized that the introduction of Bluelight should not slow down the regular treatment process.

6.5 The Patient Ombudsmen

Although we used the interview guide as described in Section 5.3, these interviews turned out a bit different. The patient ombudsmen were very talkative, and as a result we got more discussions going. Because of their type of work, they had a different approach towards the PHR than other health care actors, and many interesting opinions. We therefore decided to analyse these interviews separately.

The first section deals with meanings about the health care system today. The second section deals with meanings about the PHR in general, while the third section is the most relevant regarding our interview goals since it contains feedback on our access control model.

6.5.1 The Situation Today

One of the first things that came up during the interview with Nils Hybertsen, patient ombudsman in Sør-Trøndelag (ST), revolved around communication flow in the health care sector. The young patients who had grown up in our modern society expected that this was already well developed like in other governmental and private institutions. This was, however, not the case, and it could be a great source of dissatisfaction:

> **Interviewee:** *And when it comes to a potential producer of dissatisfaction with the health care service, this is a ticking time bomb.*

Over 50 percent of patient inquiries to the patient ombudsman in ST today were related to dissatisfaction with communication. The most critical point was communication flow from one level of care to another (e.g. hospital physician to home care services). This was often where things went wrong. Kjell J. Vang, patient ombudsman in Nord-Trøndelag (NT), also pointed

out that there existed responsibility problems in the health care sector. The patient did not know who to deal with, and the health care actors did not know about each other and what services each actor provided. A lot of improvement could according to him be done on this field.

6.5.2 General Thoughts About the PHR

Both of the patient ombudsmen acknowledged that a PHR could be useful, especially for many patients with chronic illness. But not everyone would benefit from such a solution. Some might not be able to handle a PHR, as pointed out by the patient ombudsman in ST:

> **Interviewee:** *Healthy people in a quite normal life situation seems OK, they can for instance choose that those shall get access, and those shall not. But others would get a headache as early as* (...)
> **Interviewer:** *Yes. Because that's what has been ...*
> **Interviewee:** *Press 1 do they say on the phone, right? Then somebody hangs up.*

The patient ombudsman in NT mentioned the technical barriers that many old people had when trying to use electronic equipment. This problem would, however, decrease as newer generations with more knowledge about computers would grow up. Another problem that was mentioned, was related to consent competence. The patient ombudsman in NT meant that it would in some cases be problematic to decide who were fit to manage their own PHR, e.g. many demented people could be fully capable of taking decisions on their own. The patient ombudsman in ST warned against making the PHR mandatory for everyone because of these current problems. It could, however, be an option for those who were capable and motivated to use it.

A possible problem when giving a patient free access to his or her record could be that the patient would misunderstand what was written. The patient ombudsman in ST reported incidents where this had happened with the paper based record.

One of the key functionalities with the PHR is to, as a patient, be able to contribute to the record and communicate with health care personnel. The patient ombudsman in NT meant that there were both negative and positive issues regarding this. On the positive side, the patient could to a larger extent be able to administer the content in the record. On the negative side, there was a risk that the patient would fill the PHR with irrelevant information that would make it harder for health care personnel to get hold of the relevant parts. In addition, by allowing the patient to communicate with for instance his or her primary physician, one would risk to loose a part of the physical contact with the health care service. This would be problematic, as explained by patient ombudsman in NT:

> **Interviewee:** *Because there are many possibilities for a patient to deliberately or unconsciously manipulate personal information. Right? Make the situation appear more serious than it is. While at the physicians office you are to a larger degree observed. It is very difficult to observe a patient via the internet.*

When discussing security and safety, the patient ombudsman in NT was very clear that the PHR should not worsen this. The patient ombudsman in ST also said that the PHR should not complicate the information flow between health care actors.

6.5.3 Our Access Control Model

The patient ombudsman in NT thought that relevant roles in the PHR could be boundless, as long as they were health related. Based on his experiences with patients who had contacted him, it was important that they should have an opportunity to control the information in their record, and decide what information to pass on from one health care actor to another. Except from when essential information would be needed in order to provide proper treatment. The patient ombudsman in NT saw the possibility that some patients might choose to retain this essential information. The patient ombudsman in ST thought the access control model seemed like a good solution for today, but emphasised that in the future, things might change and another type of model could be more applicable.

CHAPTER 7

Access Control Model

This chapter presents the final access control model for the personal health record (PHR), developed during this thesis.

7.1 Changes Made From the Tentative Access Control Model

This section points out issues that were raised in the autumn project "Access Control Model for Personal Health Record" [1] and during the analysis of the new interviews. It is meant to clarify any changes made to the tentative access control model from the autumn project.

In the autumn project we initially suggested a list of roles, institutions and resources for the PHR. The development of the resource list was based on existing literature and standards. The pilot interviews gave little new information on this subject, and it also turned out to be too time consuming. We have therefore decided not to work further with this list. Nor will our institution list be further developed, due to the same reasons, but instead it will work as a foundation for a conceptual institution hierarchy. The original lists can be found in the autumn project [1].

Our initial role list was also quite substantial, based on literature, standards and pilot interviews. We have, however, decided to narrow the list down to only some key roles that we think will be relevant when first introducing the PHR. These roles are based on the original role list in addition to input from the interviews. Further expansion of this list to suit the use will be necessary before the PHR is deployed.

7.1.1 Issues Noted in the Autumn Project

The following sections contain issues we brought up in the discussion of the autumn project.

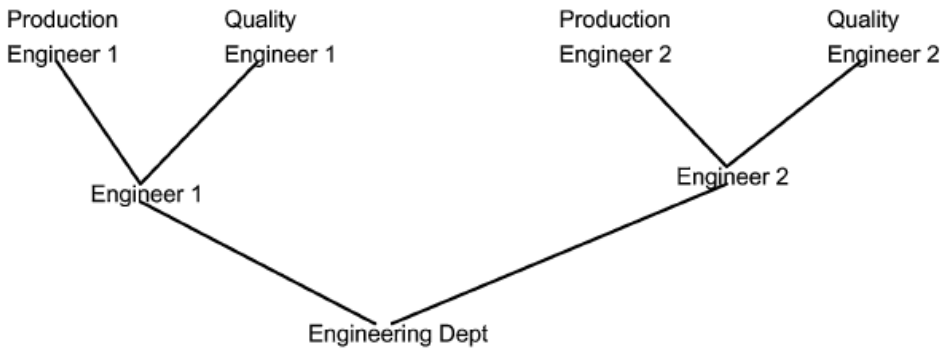

Figure 7.1: Example role hierarchy. Taken from [20]

Hierarchy

After the evaluation of the pilot interviews and the tentative access control model from the autumn project, we thought of the possibility of using a role and institution hierarchy. The main idea is that there exist specifications of both roles and institutions that will easily fit in a hierarchy. Ferraiolo et al. [20] describes role hierarchy as "an inheritance relation among roles", so that r_1 inherits role r_2 if all privileges of r_2 are also privileges of r_1. This way the hierarchy can be represented by a three structure, where membership is inherited top-down, and permissions are inherited bottom-up. Figure 7.1 shows an example role hierarchy.

The role hierarchy will split the different roles and categorise them in the hierarchy, letting the patient select between general or more specific roles when setting access. As we shall see, the list we have developed for our access control model contains only a limited set of roles, and it is therefore difficult to build a large hierarchical structure. However, we believe that the role list should be extended in the future, and that this will make the hierarchical structure more clear and useful. The hierarchy is built based on the original list from the autumn project and the new interviews.

The institution hierarchy will work the same way as the role hierarchy. An institution hierarchy will enable the patient the opportunity to set access on a general basis or to be more specific. The hierarchy can be built by using the organisational units of the institutions, also making it easier for the patient to understand. The institutional hierarchy is inspired by the institutions suggested in the autumn project. We will in this report only present a principal hierarchy with the different types of institutions and their position, not concrete institutions.

Fragment

According to KITH standard [24] the record should be divided into cases, documents and fragments. Our interviews indicated that splitting up information to this extent could lead to key information getting left out when setting access. We have therefore decided to not include fragment in the final version of the access control model.

Bluelight

In the autumn project we discussed the possibility of adding the Bluelight function to our access control model. We have reached the conclusion that it should be part of the overall functionality in the PHR. It is, however, not part of the access control model itself, but rather an implementation issue.

7.1.2 New Issues

In this section we look at issues that have been raised during this thesis, including the new interviews.

Record Defined Groups

The interviews led us to the idea of introducing user specified groups. This was not directly expressed by the interviewees, but we believe that it could make access control easier. Record defined groups will enable the patient to assign users to user defined groups, and further give access to the whole group rather than individual users. This can be effective when assigning access to family members, the work group in an IP, or other key persons in a patient's treatment situation.

Redefinition of Global and Record User Assignment

The *global user assignment* (GUA) and *record user assignment* (RUA) have been redefined to be between the user and the institution-role assignment instead of between the user and the roles entity. This solves the problem that one can not trace a user to an institution and back, letting the access control model contain all the information about the relationship between users, roles and institutions.

7.2 Complete Access Control Model

This section presents the complete access control model.

7.2.1 Fundamental Principles

The following principles defined in the autumn project still apply, and are the basis for our access control model:

1. The PHR is an addition to the existing systems, aiming at making information sharing between these systems easier.

We do not believe the PHR can replace the already existing systems throughout the health and social sector. However, if the PHR is just to be another of these systems the thought of the patient controlled record is not valid. It is therefore important that the PHR system is closely connected to the existing systems, making it possible to exchange information and access preferences.

2. The user must be able to understand the terminology used to control the access to his or her record, in this case the use of a subject, a resource and an operation.

The variation of patients that are potential users of the PHR is large, and emphasises the need for usability. The basics of the access control model has to be easy to understand, and is based on the three key words subject, resource and operation. The subject is the target user, role or institution of the access control, the resource is the information to be controlled, and the operation defines the access rights for the previous.

3. The user must be able to adjust access with a high level of granularity, in our case from institutions and cases down to levels of individual users and documents.

For the PHR to be as flexible as possible it is important to have a high level of granularity when setting the access control. In our model this has been solved by letting the user choose to set access to users, roles, institutions, or groups, and to divide resources into cases and documents. It is, however, important that the degree of flexibility does not make it to complicated for the user.

7.2.2 Subject, Resource and Operation

As already mentioned, subject, resource and operation is the core of the access control model. The three key words define who should have access to what.

Subject

The subject represents the target of the access control. This can in our model be a specific user, a role, an institution, or a combination of these. We have also in our complete model introduced the possibility for the user to define their own user groups, and allowing access to these groups in the same way as the others.

The role hierarchy can be seen in Figure 7.2. As mentioned, it is based on the original role list from the autumn project in addition to the new interviews, and shows the key roles we have reached in discussion with Jorunn Bjerkan. Further hierarchical expansion of the different roles, e.g. nurse, can be done according to the needs.

The principal institutional hierarchy is shown in Figure 7.3. It shows the division of the organizational units, and gives an impression of the complexity involved. In an implementation of the PHR, this hierarchy has to be built with the real institutions that are needed. Note that this only shows some of the different institutions that could be a part of the PHR.

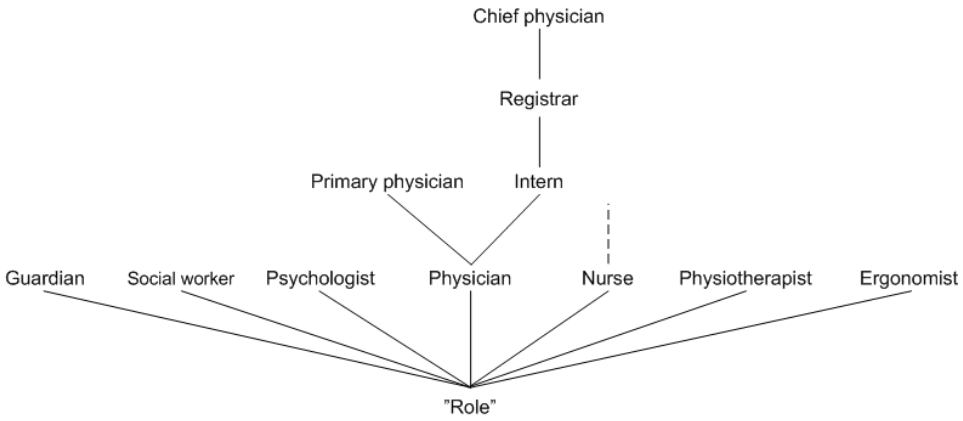

Figure 7.2: Role hierarchy

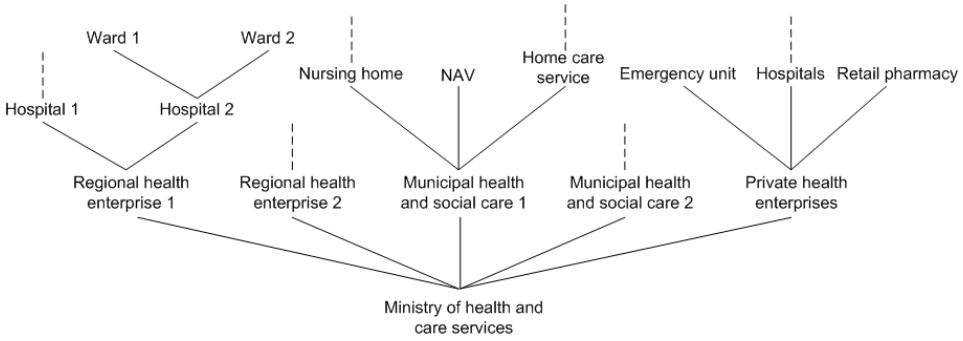

Figure 7.3: Institution hierarchy structure

Resource

Resources represent the information that is being controlled by the access control. Figure 7.4 shows that resources are a generalisation of cases and documents, where cases can be a collection of documents and possibly other cases. By letting the patient set access on both case-level and document-level gives a lot of opportunities.

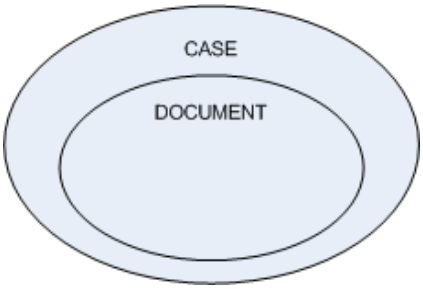

Figure 7.4: Resources

No further work on specific resources has been done in this thesis, but a suggestive list was made in the autumn project, as mentioned earlier.

Operation

The operations define what the user can do to a resource. From the autumn project we have the three standard operations:

1. No access
2. Read access
3. Read and write access

In addition to this the access control need mechanisms to limit only health care actors to write medical documents.

Our access control model differs from the RBAC-model in that one operation actually consists of a set of legal operations. The RBAC-model has a relation between each allowed operation in the system and the resources. However, to be able to deny access in our model we have had to change this so that an operation rather has to be regarded as a set of operations that are allowed. This creates access levels that can be represented by an integer, where in one end you have full right to read and write, and in the other end you have a deny functionality. One would thus assign access with a single integer rather than a list of operations.

7.2.3 Access Control Model Formalism

This section presents the definition of the access control model we have developed during this thesis. It is a continuance of the tentative access control model developed in the autumn project[1] with changes as described in the previous section.

[1]Hence the reuse of the description with some changes.

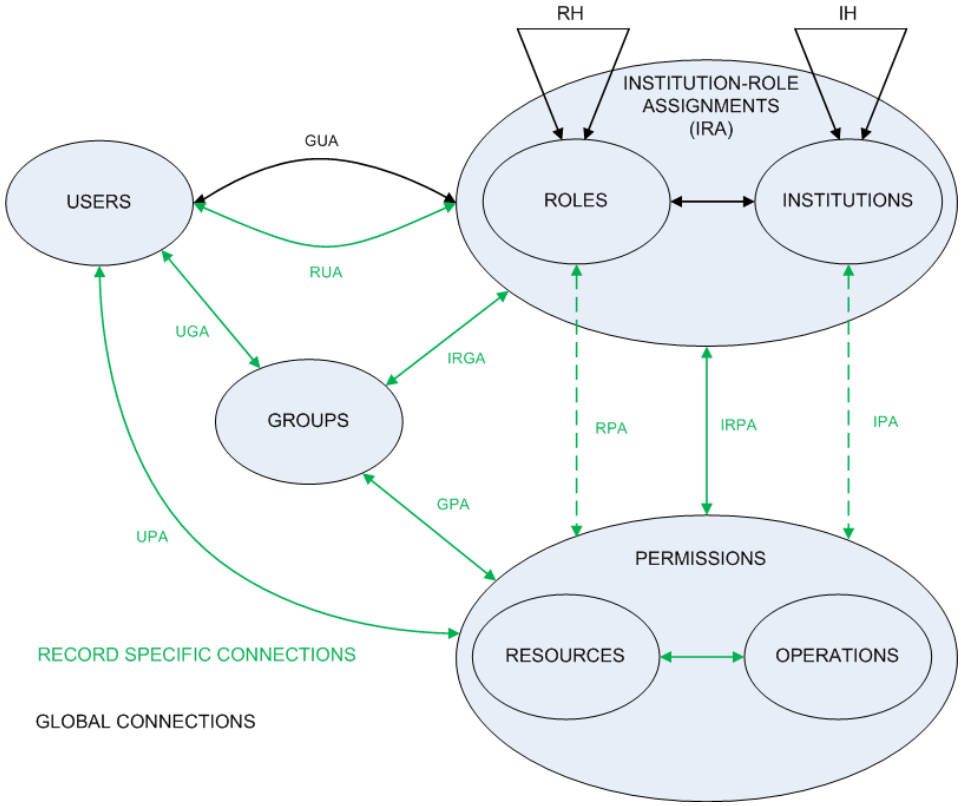

Figure 7.5: Access control model for PHR. Global connections are marked in black and record specific connections in green

The model is depicted in Figure 7.5, and shows the data elements and relations that defines the model. It is important to understand that in the model, access is given to resources in one specific record, and not to resources in general. This is shown in the figure as two different colours, global connections in black and record specific connections in green. The data elements consist of USERS, ROLES, INSTITUTIONS, GROUPS, RESOURCES, OPERATIONS, and PERMISSIONS. The first four data elements correspond to the subject part, while the rest corresponds to resources and operations, further described in Section 7.2.2.

The model also illustrates the relations depicted as arrows between the data elements. PERMISSIONS are the result of the relation between RESOURCES and OPERATIONS. Access is then given based on rules formed by permissions given to a subject. Figure 7.5 shows the many-to-many relation between resources and operations which form the permission that denotes an approval or denial to perform an operation on one or more resources.

The many-to-many relation between roles and institutions, meaning that a role can be present in many institutions, and that one institution can have many different roles assigned, forms the relation INSTITUTION-ROLE ASSIGNMENTS (IRA). This relation is used when assigning access for roles and institutions. It is important to notice that a role not necessarily needs to be designated an institution.

The *global user assignment* (GUA) and the *record user assignment* (RUA) relations both show that a user is assigned one or more combinations of roles and institutions through IRA, and a institution-role combination can be assigned to one or more users. While the GUA applies to the whole system, the RUA only applies to one specific record, shown by the colour difference seen in Figure 7.5. In this way, roles that are specific for each patient, e.g. a primary physician, can be mapped to a user in each record using RUA, while a user can be mapped to the role physician throughout the whole system using GUA.

User group assignment (UGA) and *institution-role group assignment* (IRGA) are record specific relations respectively assigning users and institution-role combinations to a group.

Both roles and institutions have a hierarchical structure. This is shown in the figure by the *role hierarchy* (RH) and the *institution hierarchy* (IH) relations. This means that a role or institution can inherit membership from another role or institution, and that way also inherit its permissions. See Section 7.1.1 for more information.

In Figure 7.5 we can see the permission assignment relations between the subject, top, and the permissions, bottom. The *user permission assignment* (UPA) is used when the subject is a user; access is then given to a specific user. To be able to assign permissions to roles in general, and to roles affiliated a specific institution we have introduced the *institution role permission assignment* (IRPA). IRPA is a combination of a role permission assignment (RPA) and an *institution permission assignment* (IPA) denoted by dotted lines [2]. RPA gives a role at any institution permission, while IPA gives an institution permission. Because we do not want to give everybody at the institution the same access, but rather assign different permissions to the roles affiliated to the institution, the relation between roles and institutions have to be used. The model solves this through the use of the IRPA relations which utilise IRA. By using wildcards in IRPA we can also give access to all roles at an institution, or all roles disregarding affiliations to institutions. *Group permission assignment* (GPA) is used for setting permissions to record specific groups, consisting of user and institution-role combinations.

[2]Note that RPA and IPA are only part of the model for explanatory purposes.

We can therefore denote the following rules for defining access:

- `UPA: (P, U) where P is a permission and U a user.`
- `IRPA: (P, (R, I)) where P is a permission, R a role and I an institution.`
- `By setting R as wildcard _ we get (P, (_, I)) which denote that P is assigned to all roles affiliated to institution I`
- `By setting I as wildcard _ we get (P, (R, _)) which denote that P is assigned to all roles disregarding affiliation to an institution`
- `GPA: (P, G) where P is a permission and G a group.`

This ability to give permissions to both users and institutions in addition to roles, provides great flexibility and granularity. It will enable the owner to adjust access from strict, giving permissions to named users, to a more open and general access using roles and institutions. Introducing record specific groups that can be granted permission, in addition to the division of resources into cases and documents, further enhance the flexibility of the access control model.

The following is a formal summary of our access control model:

- *USERS, GROUPS, ROLES, INSTITUTIONS, RESOURCES,* and *OPERATIONS*, basic data elements.
- $IRA \subseteq ROLES\ x\ INSTITUTIONS$, a many-to-many mapping role-to-institution assignment relation for the whole system.
- $GUA \subseteq USERS\ x\ IRA$, a many-to-many mapping user-to-institution-role assignment relation for the whole system.
- $RUA \subseteq USERS\ x\ IRA$, a many-to-many mapping user-to-institution-role assignment relation for a specific record.
- $UGA \subseteq USERS\ x\ GROUPS$, a many-to-many mapping user-to-group assignment relation for a specific record.
- $assigned\_roles(i{:}institution) = \{r \in ROLES \mid (r,i) \in IRA\}$, mapping of institution i onto a set of roles.
- $assigned\_users\_global(q{:}IRA) = \{u \in USERS \mid (u,q) \in GUA\}$, global system mapping of an institution-role assignment q onto a set of users.
- $assigned\_users\_record(r{:}ROLES) = \{u \in USERS \mid (u,r) \in RUA\}$, record mapping of role r onto a set of users.
- $assigned\_user\_groups(g{:}GROUPS) = \{u \in USERS \mid (u,g) \in UGA\}$, record mapping of group g onto a set of users.
- $PERMISSIONS = 2^{(OPERATIONS\ x\ RESOURCES)}$, the set of permissions.
- $UPA \subseteq PERMISSIONS\ x\ USERS$, a many-to-many mapping permission-to-user assignment relation.

- *IRPA* $\subseteq$ *PERMISSIONS x IRA*, a many-to-many mapping permission-to-institution-role assignment relation.
- *GPA* $\subseteq$ *PERMISSIONS x GROUPS*, a many-to-many mapping permission-to-group assignment relation.
- *assigned_permissions(u:USERS)* = $\{p \in$ *PERMISSIONS* $|$ *(p,u)* $\in$ *UPA*$\}$, the mapping of user u onto a set of permissions.
- *assigned_permissions(q:IRA)* = $\{p \in$ *PERMISSIONS* $|$ *(p,q)* $\in$ *IRPA*$\}$, the mapping of an institution-role assignment q onto a set of permissions.
- *Ob(p:PERMISSIONS)* $\rightarrow \{$*op* $\subseteq$ *OPERATIONS*$\}$, the permission-to-operation mapping giving the set of operations associated with permission p.
- *Ob(p:PERMISSIONS)* $\rightarrow \{$*re* $\subseteq$ *RESOURCES*$\}$, the permission-to-resource mapping giving the set of resources associated with permission p.
- *RH* $\subseteq$ *ROLES x ROLES* is the role inheritance relation, written as $\succeq_r$, where $r_1 \succeq_r r_2$ only if all permissions of r_2 are also permissions of r_1, and all users of r_1 are also users of r_2. Formally: $r_1 \succeq_r r_2 \Rightarrow$ *authorized_permissions*$(r_2,*) \subseteq$ *authorized_permissions*$(r_1,*) \wedge$ *authorized_users*$(r_1,*) \subseteq$ *authorized_users*$(r_2,*)$, where * is a wildcard.
- *IH* $\subseteq$ *INSTITUTION x INSTITUTION* is the institution inheritance relation, written as $\succeq_i$, defined the same way as RH. Formally: $i_1 \succeq_i i_2 \Rightarrow$ *authorized_permissions*$(*,i_2) \subseteq$ *authorized_permissions*$(*,i_1) \wedge$ *authorized_users*$(*,i_1) \subseteq$ *authorized_users*$(*,i_2)$, where * is a wildcard.
- *authorized_users(r:ROLES, i:INSTITUTION)* = $\{u \in$ *USERS* $|$ $r' \succeq_r r$, $i' \succeq_i i$, *(u,r´,i´)* $\subseteq$ *GUA* $\cup$ *RUA*$\}$, the mapping of role r and institution i onto a set of users in the presence of a role and institution hierarchy.
- *authorized_permissions(r:ROLES, i:INSTITUTION)* = $\{p \in$ *PERMISSIONS* $|$ $r' \succeq_r r$, $i' \succeq_i i$, *(p,r´,i´)* $\subseteq$ *IRPA*$\}$, the mapping of role r and institution i onto a set of permissions in the presence of a role and institution hierarchy.

Appendix C shows a practical example of how the relations and assignments can be set up in a record using the described access control model.

7.2.4 Precedence Rules

The order of which the permission assignments are evaluated in the system can change the outcome, and a clarification is therefore necessary. There are three dimensions that afflict the evaluation of access:

Subject: order of evaluation of user, group and institution-role assignment.

Resource: order of evaluation of document and case.

Operation: order of evaluation of no access, read access, and read and write access.

By changing the order within these dimensions in addition to the order amongst them, the evaluation of a permission assignment will change, and it is therefore important for our access control model that rules for evaluation are stated.

1. The precedence of the operations for users should be: no access, read and write access, read access. For groups and institution-role assignments it should be: read and write access, read access, and no access.

This means that a rule based on a user evaluated to no access should always take precedence before a read and write access, which again should take precedence before a read access. For groups and institution-role assignments the permission giving the most access should apply.

2. The order of evaluation of the subject should be: user, group, institution-role assignment.
   (a) In the case of evaluating the hierarchies for the latter two, this should be done in a breadth-first search.
   (b) The precedence of operations for the parents should be according to rule 1.

This means that a rule regarding a specific user should always take precedence before a group, which again should take precedence before an institution-role assignment.

3. The order of evaluation of the resource should be: document, case.
   (a) If a resource is not found, its parent cases should be evaluated, one level at the time in a breadth-first search.
   (b) The precedence of operations for the parents should be: read and write access, read access, and no access.

This means that a rule regarding a document should always take precedence before a case. Further if the resource has no matching policy, the resource should inherit one of its parent's permissions.

4. The order of evaluation of a set of permission assignments should be subject, resource, operation.

This means that when comparing a set of possible permission assignments one should always choose the permission assignment that first matches the subject with highest precedence, then the resource with highest precedence, and finally the operation with highest precedence, all according to rules 1-3.

Having the tuple (S,R,O) where S is the subject, R the resource and O the operation, we have the following examples:

- `(user, document, read and write access) will take precedence over (group, document, read access)`

- (user A, document, no access) will take precedence over (user A, document, read and write access)
- (group A, document B, read access) will take precedence over (group A, document B, no access)
- (group A, document B, read access) will take precedence over (group A, case C, no access), if document B is part of case C

CHAPTER 8

Implementation

This chapter explains how the access control model presented in Chapter 7 can be implemented in the already existing PHR Indivo application described in Chapter 3. It explains the implementation on a high level, outlining how and where the major changes have to be done in addition to how the XACML language has to be used, without any definite code or pseudo-code.

Due to the degree of complexity in our access control model, we have looked into two possible solutions that can be implemented using the basic features in XACML. We have further looked at the possibility of a third solution we believe can be implemented using more advanced features of the XACML language.

8.1 General Challenges and Solutions

This section describes the general challenges that have to be solved for all the solutions.

8.1.1 Indivo Application

Presently the Indivo client does not support the fine-grained access control we have suggested. Therefore some major changes are needed, both on making new policies, and creating the requests. This should be done by extracting the policy set from the record, adding changes where needed on the client side, and then sending it back to the server and store.

8.1.2 Hierarchies

The major challenge in implementing our access control model is the extended use of hierarchies. Both roles and institutions are represented with hierarchies. In a request based on

these subjects, not only attributes belonging to the request have to be checked, but also other possible values according to the specific hierarchies. The role and institution hierarchy should be represented in a dynamic way, making it easy to change. Our suggestion is therefore to represent the hierarchies in instances of Java classes which can be read when needed, either by the client or the server. These Java classes can easily be created at compilation time by translating XML schema files as described in Section 3.2.3.

The resource hierarchy also constitutes a challenge. When a request is sent to be evaluated by the RecordPDP with a single specific resource, this resource may not be present in any of the policies. However, the resource, a document or a case, can be part of another case, and ergo inherit its permissions. Further request for access to the parent resources will therefore have to be done to reach a conclusion. Evaluation of the parent resources should be based on the rules in Section 7.2.4.

To manage relations between cases and documents we suggest the introduction of a fourth core document in the record, a RelationDocument. This document, in addition to the resources themselves, can contain the relation information, making it structured and easy for lookup. Another student project at NTNU has come up with a suggestive solution for this relation document [25].

8.1.3 Record Specific Assignments

In the access control model we have defined the possibility for the patient to define their own groups in addition to record specific roles. As this information should be located in the requested record, and not in the requester's record, it can not be a part of the original request. To solve this we suggest the introduction of a fifth core document in the Indivo record structure. This document, RecordSpecificAssignmentDocument, should contain a mapping of users to groups (UGA), institution-role combinations to groups (IRGA), and users to record specific roles (RUA). The document can, in the same way as some of the other core documents, be implemented using a combination of Java and XML to define the document structure. This way we can extract the groups and record specific roles the requester is allocated to in the record, adding it to the request before sending it to the PDP. Note that we abandon Indivo's original intended group concept.

8.1.4 Role Institution Relation

As defined in the access control model, a user can be a member of multiple institution-role combinations. To maintain the relationship between the role and institution, and not loose this in the communication between the server and client, we suggest a concatenation of these two. Role and institution can be represented by a single string where a key character split the role and institution, e.g. <role:institution>. This will change the current solution in Indivo by not using a default role. Information about the user's institution-role affiliation can, as today, be stored in Indivo's ActorAttributes document. This way the user logs on with the username and password, and receives additional information about role and institution from the user's own record.

The concatenation of role and institution, however, leads to a new problem. Policies that are defined to use wildcards (_) on either roles or institutions will not get the desired effect, e.g. a

request with <physician:hospital> will not match the policy <physician:_>. To solve this, the request also have to contain each role and institution with wildcards, e.g. the above example will lead to a request also containing <physician:_> and <_:hospital>.

8.1.5 Evaluation

The evaluation of the request against the policies has to follow the precedence rules as described in Section 7.2.4. To get this desired sequence of evaluation we suggest a fixed setup for the record's policy set. First we wish to sort the policies on the resource in question, then the different actions for this resource, and finally the subject affected by the policy:

```
<PolicySet>
(This is the outer policy set in the access policy document in the record.)
  <PolicySet>
   Target: Resource
   (This policy set targets the resource in the access policy.
   One policy set for each resource.)
     .
     .
     .
    <PolicySet>
     Target: Action
     (This policy set targets the actions in the access policy.
     One policy set for each action.)
     .
     .
     .
      <Policy>
       Target: Subject
       (This policy targets the subject in the access policy.
       One policy for each subject, where subjects with usernames come
       first, then groups and finally institution-role.)
```

To achieve the right evaluation sequence of the subjects, the different policies regarding subjects also have to be ordered, i.e. policies regarding users first, then groups, and finally institution-roles. This in combination with a customised XACML combining algorithm will give the desired precedence effect. The combination algorithm that has to be implemented should have a similar effect as the First-applicable combining algorithm [11]. The First-applicable algorithm will stop at the first possible permit during an evaluation and return that decision. The customised algorithm should, however, evaluate all subjects at the same level before it returns an effect following the precedence rules. E.g. if a request contain two group subjects where the first evaluate to deny, and the second to permit, the overall effect should be permit.

8.1.6 Operation - Action

We have in our model defined some basic operations that can be given. Indivo uses a set of different actions much like these operations, and a mapping between these are shown in Table

8.1.

Our model	Indivo
No access	Deny on all actions
Read access	Permit on read actions
Read and write access	Permit on read actions Permit on update actions

Table 8.1: Operation mappings

In addition to these, Indivo also has actions for creating a new record, authenticating and sending messages. None of these are directly used in our access control model. Creating a new record is typically something that should be reserved for system administrators, authenticating is an action used by Indivo in the authenticating process, and finally sending messages is an extra feature that can utilise the access control to manage access to messages.

8.2 Solution 1

This section presents the first implementation solution of our access control model.

8.2.1 Overview

This solution solves the hierarchy problem by defining all possible policies at creation time. This means that if any of the subjects in the policy are roles and/or institutions, not only the original policy is created, but also policies for all possible combinations of roles and/or institutions according to the hierarchies. The result is multiple policies, covering all possible solutions of roles and institutions that should have the same access as the original role and institution.

This way a request containing information about the user, group, role and institution can be sent to the PDP. Here it is evaluated against the policies defined in the record as usual.

Policies created as a result of other policies should be treated as any normal policy, open for modification or cancellation. It is also important that the patient is made aware of the consequences of any such indirect policies.

8.2.2 Creating Access Policies

As already mentioned, the Indivo client presently does not support fine-grained access policies, and therefore requires large modifications to be functional. However, the code that creates the policies shall be able to create not only one policy, but multiple policies based on hierarchies. To do this it has to utilise the hierarchy structures as mentioned earlier. Note that functions for policy creation are present in the application, it is therefore only a matter of setting the right input.

8.2.3 Creating Requests

The process of creating the request has to be slightly altered. In addition to the user id already present, we also want to add information about group membership and institution-role affiliations. This should be done by querying the RecordSpecificAssignmentDocument for group membership based on the requester's username and institution-role-affiliation, in addition to any possible record specific roles. These queries can be done before the request is made in the checkPermissionAndAudit method in the DefaultActionResponder class. Further the creation of the request itself, in the RequestCtxGenerator class, has to be modified to use the group and the institution-role attribute with the wildcard.

8.3 Solution 2

This section presents the second implementation solution of our access control model.

8.3.1 Overview

Solution 2 solves the hierarchy problem the other way around. Instead of creating multiple policies, only one policy is created with the intended role and/or institution. Further the request not only contains the user's role and/or institution, but also the roles and institutions that he/she is a part of according to the respective hierarchies. This way all the possible subject values for the requester is sent in the request and evaluated against the policies in the record.

8.3.2 Creating Access Policies

Solution 2 does not require any large modifications beside the improvement of the GUI and its functionality, as only one policy is created.

8.3.3 Creating Requests

Creating requests in solution 2 is mostly similar to solution 1. However, the request also has to contain all hierarchical possibilities of roles and institutions. These additional subjects should be extracted from the files that contain the hierarchies. The main methods are as in solution 1 checkPermissionAndAudit, DefaultActionResponder and the RequestCtxGenerator class.

8.4 Solution 3

Both the previous solutions explained solve the hierarchies by adding extra information to policies or requests. This creates an unnecessary amount of information. This third solution aims at limiting the need for extra policies or requests to cover the hierarchies, and instead utilises the possibilities in XACML. The basic problem is to involve the hierarchical data structures into the evaluation of a request.

This problem can be solved by the use of custom made functions that are passed on to finder modules [15]. These can be implemented in Java, and the finder modules are sent as parameters to the PDP when it is initialised. The finder modules are therefore up to date at all times, as the RecordPDP is created at request time. By programming specific functions and finder modules we can obtain the desired evaluation process according to the access control model. The requests do not need any modifications, while the policies must have a reference to the correct custom made function. It is the functions that contain the logic to traverse the hierarchies and provide the complete set of subject attributes in the policy.

A similar solution can be used to traverse the resource hierarchy. By using a customised resource finder module, a request containing a single resource can locate other resources that have the original resource as the root of a hierarchy. As long as there is some simple hierarchy structure, the request can this way access more than one resource.

By fully using the opportunities that lies in XACML we end up with an elegant solution that does not require a lot of coding outside the XACML environment in the application, and that does not create a lot of redundant requests and policies.

CHAPTER 9

Discussion

In this chapter we discuss the problems and solutions we encountered during our work researching and developing an access control model for the PHR. Much of the foundation for our model was based on interviews with potential users of a PHR system, and we start by discussing issues regarding this process. Further discussions and comments on the access control model itself and the suggested implementation of the model in Indivo follows.

9.1 The Interviews

In order to get suggestions for important PHR roles and feedback on our access control model, we decided to arrange further qualitative interviews with patients, patient guardians and health care actors. This section contains a discussion of the results from the analysis and the interview process itself.

9.1.1 Discussion of the Analysis

The discussion of the results from the analysis is separated in three parts. First the general results are presented, which mainly consist of thoughts around the PHR and its possible future in the Norwegian health care system. The discussion of the roles and the access control model follows thereafter.

General Results

The PHR is a complex system due to several factors. The health care system is very extensive, and when including other relevant institutions like the social services it gets even more complicated. Another factor is the patient. There are as many potentially different patients

in Norway as there are inhabitants. Young and old, well and ill, and each with his or her own ideas and wishes. This puts great demands on a system that must satisfy both parties. During this thesis we have tried to solve several issues, but new ones keep emerging.

The patient ombudsmen are the most sceptical ones regarding the PHR. Among other factors they focus on security and safety, and based on work experience, they believe that many potential users do not hold enough knowledge about these topics. During our interviews we have received the same impression. The results from questions regarding computer usage, combined with security related topics that have emerged during the interviews, lead us to believe that there must be many patients and health care actors that have an insufficient and naive approach towards these topics. These are, however, barriers that hopefully will diminish as younger generations with increased knowledge about computers grow up. There are also measures that can be taken to increase knowledge and security awareness with today's generation, e.g. courses and information. In addition, several of the interviewees point out that the PHR has to be secure against malicious access. We believe that both these issues have to be one of the top priorities when developing a PHR. It contains information that is not only sensitive during a limited period of time, but as long as the patient lives and potentially also after his or her death. This focus on security and safety should, however, not affect the user friendliness of the solution in a negative way.

A PHR puts demands on a patient. He or she must actively participate, e.g. grant and deny access to parts of the record. The patient ombudsmen believe that not all patients are fit to handle such a responsibility. Some might not see the point of having a PHR and some might be too ill to manage it. Others do not have the necessary knowledge about computers and the above mentioned safety and security risks. We agree with the patient ombudsmen, and think that the PHR should not be mandatory, but an option for those who want to use it.

Another interesting topic is how actively the patient should be allowed to participate in the PHR. The patient ombudsmen and some of the health care actors are concerned about loosing the personal contact because of communication through the PHR, and that the patient provided text in the record might be meaningless and makes it more difficult to find relevant information. Both situations are unfortunate, and should be avoided. As described by one of the patient ombudsmen in Section 6.5.2, it is difficult to diagnose a patient without a proper face-to-face consultation. In addition, if the record were to fill up with unnecessary text, the health care actors might experience poorer efficiency in their work. However, some of our interviewees do not believe these problems will occur. In any case, there are ways to deal with these possible problems. One solution could be to limit where and how much the patient was allowed to write. The patient should also be made aware that the PHR is not a substitute for personal consultations with health care actors.

A very interesting find is that both patients and health care actors point out communication problems in the health care sector. Both between patient and health care actors and between different levels of health care, e.g. between a hospital and the home care services. A majority of our interviewees think that the PHR is a good idea, and one of our patients had very good experiences using the Sampro IP tool, which is similar to the PHR. These points make us believe that the PHR should be implemented in the future, despite the challenges mentioned above.

Roles

A general impression when asking the patients and patient guardians about roles, was that their knowledge on the subject was quite limited. Their answers were vague and characterised by uncertainty, and they only mentioned roles they were in contact with. These are the same experiences as we encountered with the pilot interviews in our autumn project. An example of their lack of knowledge was the use of the role physician. Some of the patients used this superior term on several types of physicians, without specifying whether it e.g. was a primary physician or specialist at a hospital. In some situations the context made clear what type it was, and in other situations we asked follow up questions in order to get the specific role. The health care actors had generally a better overview of the health care roles, even though they mainly only mentioned roles related to the type of patients they treated. This focus on relevant roles causes us to believe that our results shown in Section 6.3 are biased, since our interviewees belong to a limited occupational part of the health care sector. Our selected main roles shown in Figure 7.2 are therefore only the most general.

The Access Control Model

One of the goals in our interviews was to get feedback on our access control model. To do this we gave a basic explanation of the access control model without any technical details. The responses on the explanation and the following questions were mostly positive, accepting the access control model as a good solution in a PHR setting.

The different answers regarding access to users versus an institution-role combination, both receiving positive feedback, show the need for diversity in the access control model, according to the patients' different situations. The interviewees also supported the model's possibility to specifically deny some users access to information in the record. However, as pointed out by some of the interviewees themselves, this can lead to unfortunate and serious incidents which are not in the patients best interests. To avoid such happenings, much work should therefore be done to educate the users, in addition to introducing emergency functions such as Bluelight for the health care actors.

Our general opinion is that the interviews gave us little feedback on the concrete access control model, but rather inspired us and confirmed the wanted functionality. This is also reflected in the analysis in Section 6.4. A possible solution leading to a better evaluation of the access control model could have been to implement a pilot system. This pilot system could then be used in arranged workshops with the interviewees, with subsequent discussions.

9.1.2 Discussion of the Interview Process

This section contains a discussion around the execution of the qualitative interviews. There are three parts, the first dealing with experiences related to our interview guide. The second deals with experiences related to how we carried out the interviews, while the last part deals with the tools we used throughout the interviews and the analysis.

The Interview Guide

Before we started the new interview round we made some updates to our old interview guide as described in Section 5.3. After having completed the interviews, we were left with the impression that our new guide improved several of the problems we experienced with the old one.

One of the things we changed, was using more time explaining the PHR, introducing sketches and images to simplify the explanation. This seemed to improve the interviewees understanding, and they were better able to put themselves in a correct PHR context and give good answers. We also removed some of the main points from the guide. We were then left with more time to the most important points, and lack of time was not an issue in most of the interviews.

By changing some of our questions we tried to better allow for the interviewees knowledge. In addition, we built on our experiences during the interviews and adapted the questions to the interviewees if necessary. This was a successful approach, and we were able to get more information than in the pilot interviews. We also created patient cases based on PD in order to make the interviewees focus not only on their own situation. Experiences with this part of the interviews are described later.

Carrying out the Interview

When carrying out the interviews we did not experience any major problems. And even though we talked to some people living what can be characterised as difficult lives, the atmosphere was good during all the sessions. Their hospitality and effort were admirable, and even though there were some sad moments, coffee, cookies and laughter were more common. All the interviewees were warmed up with "small talk", which probably contributed to this favourable situation. In addition, we now had interview experience from the pilot interviews in the autumn project, and felt more comfortable with the interviewer role.

It is worth noting that we did not participate in the interviews together because of the limit of three people described in Section 5.1. It was Jorunn Bjerkan, one of us, and the other master's degree student. So even though we used the same guide, the interviews were performed with a few minor differences due to fact that we were two different individuals. However, we believe that this did not influence the results in a negative way. It also seemed like a good arrangement for the interviewee to only have one interviewer to deal with.

The only real problem encountered when carrying out the interviews, was during the phone interview. The connection was bad and hence it was sometimes difficult to hear the other part. A couple of times the line was broken and we had to recall the interviewee.

Tools and Equipment

We had decided to once more try to use PD as a tool to get feedback on our access control model. This time the workshop was better planned, with cards representing different parts of the system and cases created in order to get the patient to focus not only on his or her situation. The experiences were both positive and negative. When introducing the workshop

to the interviewees, the interview surroundings were in some of the interviews not suited for the type of work. Typically we would sit in a sofa round a small table filled with coffee cups and cookies. In those cases we omitted the workshop and started a normal discussion. We also found out that in some cases we could use parts of the PD material as an example of how to use the model, without involving the specific cases. The feedback from the interviewees was quite good when using this approach. We did, however, loose some of the rich discussions that would emerge when going through the cases.

A very interesting find was that even though we had created specific cases with an imaginary patient, the interviewees still tried to convert the situation to fit to their situation. If further interviews were to take place, a good idea might be to create the cases in a way that allowed the interviewee to use his or her situation when solving the tasks. In addition, the PD workshop should probably be performed as a separate task, and not as the last part in a lengthy interview. In that way the interviewee would be more concentrated and there would be better time. The ideal solution could be to let the interviewee solve the tasks using a pilot model on a computer, since this would be more realistic.

The sound and video quality of the interview was mostly very good. The transcribing went well, even though there sometimes was some mumbling that was difficult to interpret. When analysing the transcriptions we used a software tool called Nvivo [22] which was very useful. Due to the large amounts of data, it would have been difficult to get an overview of all the meaning units if we only were to use printed material. With Nvivo we were able to group meaning units together, and keep good control of all the transcriptions.

9.2 The Access Control Model

The development of an access control model for the PHR has been our main focus throughout this thesis. This section contains a discussion of this model. First we look at the general functionality of the model and what it can offer to a user of the PHR, and then some issues concerning the possible implementation of the model.

9.2.1 General Functionality

The overall object of the development of the access control model has been to find a solution for access control that puts the patients in a more central position, letting them access and control more of their information in the record. This will hopefully lead to more patient involvement, and patient centred health care. We have, based on this, tried to give the patient a variation of options for access in the model.

Our development of the model is based on RBAC, in addition to studies on access control in EHR systems and systems with similarities to the PHR, as described in the autumn project [1]. The access control model is also based on various KITH standards, something we believe is important for future interoperability. The interviews have also been a valuable source of ideas and inspirations, together with more informal discussions with experienced health care actors.

The basis for the access control model and how access is given lies in the three key groups subjects, resources and operations. Together these three can form tuples that define access in a precise way.

Subjects

The standard RBAC uses only roles as subjects. We, however, wanted a richer set of subjects according to both patient needs and the health care structure, and added basic users, groups, and institutions in addition to roles. This enables the patient to choose different levels of access, depending on the situation. We believe users, groups, roles and institutions should be enough for the patient to set the right access for subjects, however, we see the danger of patients setting access either too narrowly or too open by using only users or roles and institutions respectively.

Resources

The resources are the cases and documents, as defined by KITH. These are ordered in a hierarchy according to relations between them. By using both cases and documents the patient can choose to set a narrow access on a single document, or to give access on a wider scale through a case involving other cases and documents. As mentioned in Section 7.1.1, we have discussed the possibility of also using fragments. We have, however, decided not to include fragments as part of the resources because they probably would lead to too much fragmentation of information. By going into this level of detail, key elements could be left out making it hard for health care actors to get a hold of necessary information. Either as a result of the patient setting access to only a fragment, or by receiving access to a document missing one or more key fragments.

Operations

Our model includes three levels of access, from no access to read and write access. We believe that these basic operations are sufficient, but the set can be expanded if needed.

9.2.2 Key Features

Some of the key features of the model depicted in Figure 7.5 are discussed in the following sections.

Institution-Role Assignment

Early in the development process we decided to create a relation between roles and institutions. This relation combines a role with an institution, which is used for access. This way a patient can give access not only to a specific role, which is a fairly open access, but also to a specific role at a specific institution, which in turn is a more specified access, enhancing the granularity of the subject.

Role and Institution Hierarchy

We have introduced hierarchies for roles and institutions. In this way roles and institutions can inherit permissions and properties, and access can easily be dealt to larger or smaller number

of roles and institutions. If presented in a good way in a graphical interface, it might also help the patient be more precise in the process of setting access to different parts of the record.

Both the suggested hierarchies can be further expanded to fit the needs of the PHR and its users. However, it is important that they are not expanded beyond the comprehension of the patient as this will only give a negative effect. In general, we believe that some patients most likely will have problems understanding it, something that should be dealt with during a training phase.

Further, especially the role hierarchy can contribute to maintain the existing distribution of responsibility between the roles that exist in the health care system today, where many perform the same function, but have different roles. An institution hierarchy can also increase the patient's understanding of institutional structure and division into smaller units.

Global and Record Specific Connections

We have in our model introduced not only global assignments, but also record specific assignments. While the global assignments apply for all users, record specific assignments lets the user assign specific roles to specific users. This way, key roles that may vary from record to record can be specified, such as the role of a primary physician. We believe having such record specific roles are important to utilise roles that are personalised for each patient.

In addition to enabling record specific roles, record specific connections also enable the use of record defined groups. The individual user can this way define personalised groups consisting of users and institution-role combinations. By using these groups the patient can easier adjust access for multiple users defined by the patient self.

9.2.3 Implementation Issues

Certain features concerning the access control model can not be modelled in the figure, but should nevertheless be commented. The following sections discuss precedence rules, the possibility of Bluelight, and the creation of resources and records in the PHR.

Precedence Rules

The evaluation of a request has to follow predetermined rules as explained in Section 7.2.4. The choices made in that section have effect on how the access control model should be implemented, and if not followed, the evaluation will be different than initially intended.

We have defined that the order of evaluation of subjects should be user, group and institution-role assignment. This is based on a simple weighting of the different subjects. If an access rule targets a user, the intent for this rule is much clearer than one targeting an institution-role assignment. Further, groups are also custom specified in a record, and should therefore go before an institution-role assignment, but not before the user.

The resources are defined to be evaluated in a document-case order. Again an access rule targeting a document, should have precedence over a case because a document is more specific. Further, if the evaluation is inconclusive, the resource inherits its parent's permissions. We

have in this case decided to use a breadth-first search, searching all the parents and using a precedence of read and write access, read access, and no access. This precedence rule lets a resource with two parents, one accepting access and one denying access, inherit the accepting access, as we believe this should weigh more in such a situation. The same should apply to institution-role assignments.

The general precedence rule define that evaluation should be done on subject first, then resource and finally operation. This results in an evaluation that we believe will give the best outcome.

Bluelight

We believe that a Bluelight function should be implemented in the PHR. It should be available for the health care actors in an emergency situation. As we see it there are two options for the function. Either the patient defines who should have access to what through Bluelight, or it is predefined in the system. Due to the importance of such a function, we believe the latter option should be selected. The assignment of Bluelight to the roles physician and nurse, with the underlying roles, should be sufficient in our opinion. In addition, according to record structure, only vital parts of the record should be accessible through Bluelight.

Creation of Resources and Records

The creation of new resources and records should also be controlled. A regular patient should not be able to add all types of documents to a record, but be limited to add documents of a none-medical content such as contact information and annotations. Further the creation of a new record should exclusively be a job for system administrators.

9.3 Implementation

In Chapter 8 we presented a conceptual implementation of our access control model in Indivo. This section discusses implementation obstacles, in addition to the advantages and disadvantages of the three solutions.

9.3.1 Obstacles

The present Indivo version is at a beta state of development, and is therefore prone to changes. In addition, the present client does not support the fine grained access levels of our model, as already mentioned. Because of this, many changes have to be done to the application before our model can be fully implemented.

Since both our role and institution hierarchies are prone to change, the support for dynamical updates of these is important. By representing the hierarchies in XML-files, we mean this can be achieved. The resource hierarchy, containing the relations between the documents and cases, and the record specific assignments should both be represented in java and XML, as described in Section 8.1.3. That way they can be managed the same way as the other record specific documents.

As described in the implementation, the role and institution relation should be maintained by setting the two together in a tuple. This again creates problems where either roles or institutions are used without relation to the other. The solution to this has been to add wildcards to each in the request, resulting in three subjects. The relation is necessary, and we currently see no other way of solving this problem.

One of the larger challenges in the implementation is to get the precedence rules right in relation to the evaluation. To achieve this we need to sort the policies in a predefined order, and create a customised combining algorithm. Creating this algorithm requires good knowledge of the possibilities in XACML.

We have in the implementation done a simple mapping of the access control model's operations to the Indivo's actions. Especially the no access operation is important as it specifically denies access to a resource. It is possible that the operations in the access control model could have been more specialised, e.g. adding update and create, thus making it more similar to Indivo. We have, however, focused on the basic operations that are needed in the access control model.

9.3.2 The Solutions - Advantages and Disadvantages

Both solution 1 and 2 involve comprehensive changes in the source code. These changes might expose new and unforeseen errors in the Indivo application, and an overall desire should hence be to limit this kind of work. We believe solution 3, that instead utilises the advanced features of XACML, shows that this is possible.

When requesting a resource, solution 1 and 2 also have to form multiple requests if there are no matching policies for the original resource. This has to be done in order to traverse the resource hierarchy searching for an applicable policy. The solution is vulnerable to mistakes if the traversing process is not properly implemented. For larger systems, such a solution could also constrain the overall performance.

In general for all solutions we have the problem that for a positive evaluation of a request it is not possible to determine for what subject attribute the permission was granted. This can potentially lead to misunderstandings, and is something that should be looked further into.

All the three solutions support changes and modifications of the policies with relative ease.

Solution 1

Solution 1 is based on creating multiple policies according to the role and institution hierarchies. This creates a situation where a single request with basic information about the subject, resource and operation can easily be matched with policies. In addition, these policies, can easily be distributed to other health care systems. However, the creation of multiple policies creates an unnecessary information excess that should be avoided. The number of policies can with this solution become comprehensive. Creation of all the policies also creates a problem in relation to updates in the hierarchies. Existing policies made from previous hierarchies will not be updated, and this can therefore lead to errors in an evaluation if the hierarchies have been updated. The policies can this way be seen as static in relation to the hierarchies.

Solution 2

Solution 2 solves the previously mentioned problem with static policies by resolving the hierarchies at request time. By extracting other possible subjects and using them in the request, and only having the original policies in the policy store, the authorisation will always be up to date with changes in the hierarchies. The policies together with the request can this way be seen as dynamic in relation to the hierarchies. However, the solution requires all possible subjects to be resolved and added to the request, making the request rather large.

Solution 3

The overall advantage of solution 3 is that it does not require the comprehensive changes in the source code as the two first solutions. However, this solution requires thorough knowledge about XACML and its possibilities in order to program the necessary features. By using these advanced possibilities of the access control language, only simple and straightforward requests and policies should be necessary as both subject and resource hierarchies are resolved in the PDP.

9.4 Further Work

The qualitative interviews have shown that the access control model in general satisfies the needs of the potential users. However, there are a few points that require further work.

Both the role and institution hierarchies are far from complete. We have only been able to identify the most central roles because of our interview group with participants involved in a limited occupational part of the health care system. Participants from other parts should also be involved in order to get a better picture. Finding relevant institutions has not been part of this thesis, and should be investigated further. The same goes for relevant documents and cases that should be present in a PHR. We consider the latter task as comprehensive, and it should probably be a project on its own.

The access control model has already been tested on a general level in the qualitative interviews. A conceptual implementation is also given in this thesis. However, the model should be implemented and tested in a real system, e.g. Indivo. In that way, the potential users can get hands-on experience from using the model in a more realistic environment. The test results would probably be more thorough, and the correctness of the precedence rules in the hierarchies from a user's point of view could also be tried out. Both regarding the wishes and the security of users.

CHAPTER 10

Conclusion

This thesis presents an access control model for the personal health record, based on literature studies and qualitative interviews with potential users. The model has been developed with the patient in mind, and the health care domain has also been taken into consideration. It should therefore be well suited for the Norwegian health care system. The model supports setting access with different granularity, as well as supporting role and institution hierarchies and user defined groups. Relevant roles involved in a personal health record have also been identified through the qualitative interviews. Finally, we have determined that the Indivo personally controlled patient record can be adapted to allow for the model, and verified it with a conceptual implementation.

Of our findings, the most important ones have come as a result of the pilot interviews where we tested the model and asked about roles. By interviewing patients, health care actors and the patient ombudsmen in Sør- and Nord-Trøndelag, we have received valuable feedback. This has helped us on the path to create a complete access control model with the core roles involved. The interviewees have generally been pleased with our work, and this shows that the model seems to fulfil the wishes of both patients and health care actors. The patient ombudsmen are, however, somewhat sceptical towards introducing the personal health record as a tool for every patient. We agree on this point, and believe that it should be up to the patient whether his or her personal health record should be created or not.

Through the conceptual implementation, we have also shown that the access control model can be used in a personal health record such as Indivo, although Indivo requires substantial modifications to be able to utilise its full potential.

We consider the access control model complete and possible to implement. We do, however, believe that it should be further tested against relevant users before being taken into use. A pilot implementation on a closed server could be a good way to do this. In addition, we have only introduced the most relevant roles involved in a personal health record. This is due to time limitations and the type of interview group. Further interviews with other types of patient and

health care actors groups are required to get the full picture. Finally, the role and institution hierarchies need to be made complete.

This thesis has shown that our model with inspiration from the theory of role based access control, is a possible solution for letting patients govern their medical data. The concept of the personal health record has generally been well received in our interviews , but there are many obstacles to overcome before it can be realised. We do, however, believe that the personal health record is one of the solutions for the future in order to provide better quality of care.

APPENDIX A

Glossary

EHR:	Electronic Health Record
GPA:	Group-Permission Assignment
GUA:	Global User Assignment
Health care actor:	An employee within the health care sector, e.g. physician, nurse
IH:	Institution Hierarchy
IPA:	Institution-Permission Assignment
IRA:	Institution-Role Assignment
IRGA:	Institution-Role-Group Assignment
IRPA:	Institution-Role-Permission Assignment
JAXB:	Java Architecture for XML Binding
KITH:	Norwegian Centre for Informatics in Health and Social Care
NT:	Nord-Trøndelag
PDP:	Policy Decision Point
PEP:	Policy Enforcement Point
PHP:	Hypertext Preprocessor
PHR:	Personal Health Record
RBAC:	Role-Based Access Control
RH:	Role Hierarchy
RPA:	Role-Permission Assignment
RUA:	Record-User Assignment
SINTEF:	The Foundation for Scientific and Industrial Research at the Norwegian Institute of Technology (NTH)
SMS:	Short Message Service
ST:	Sør-Trøndelag
UGA:	User-Group Assignment
UI:	User Interface
UPA:	User-Permission Assignment
XACML:	Extensible Access Control Markup Language

Continued on next page

Abbreviation:	**Meaning:**
XML:	Extensible Markup Language

APPENDIX B

XACML Example

The following is an example of how a policy, a request and a response in XACML can look like. The example policy is taken from the Indivo system [14].

B.1 Policy

When a user sends a Create action request from the client to the Indivo server, this is the policyset which the server PDP uses to decide whether the user is allowed to perform the action. The policyset has one target and one policy. The target specifies that the policyset applies to all subjects and all resources, but is limited to actions of the type *urn:org:chip:ping:action:create*. The policy contains a rule that specifies that only subjects with the attribute id *urn:org:chip:ping:attribute:role* set to *administrator* are given permit.

```
<PolicySet
 xmlns="urn:oasis:names:tc:xacml:1.0:policy"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 PolicySetId="urn:PolicySetId"
 PolicyCombiningAlgId="urn:oasis:names:tc:xacml:1.0:policy-combining-
  algorithm:ordered-deny-overrides">
 <Description>
  This PolicySet applies to all subject attempting to perform a
  Create action through the DefaultCreateActionResponder.
  The Policies included in the set should be written to evaluate PERMIT if
  the action is allowed.  The absence of an applicable policy or a policy
  that evaluates to DENY will cause the responder to not allow the
  creation to proceed.
```

```
</Description>
<Target>
 <Subjects>
  <AnySubject/>
 </Subjects>
 <Resources>
  <AnyResource/>
 </Resources>
 <Actions>
  <Action>
   <ActionMatch MatchId="urn:oasis:names:tc:xacml:1.0:function:
    string-equal">
    <AttributeValue DataType="http://www.w3.org/2001/XMLSchema#string">
     urn:org:chip:ping:action:create</AttributeValue>
    <ActionAttributeDesignator DataType=
     "http://www.w3.org/2001/XMLSchema#string"
     AttributeId="urn:oasis:names:tc:xacml:1.0:action:action-id"/>
   </ActionMatch>
  </Action>
 </Actions>
</Target>
<Policy
 xmlns="urn:oasis:names:tc:xacml:1.0:policy"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:schemaLocation="urn:oasis:names:tc:xacml:1.0:policy
  cs-xacml-schema-policy-01.xsd"
 PolicyId="GeneratedPolicy"
 RuleCombiningAlgId="urn:oasis:names:tc:xacml:1.0:rule-combining-algorithm:
  ordered-deny-overrides">

 <Description>
  This Policy applies to the same Target as its parent PolicySet.
  The single Rule in this policy permits subjects with a role of
  "administrator" to perform a creation.
 </Description>

 <Target>
  <Subjects>
   <AnySubject/>
  </Subjects>
  <Resources>
   <AnyResource/>
  </Resources>
  <Actions>
   <AnyAction/>
  </Actions>
 </Target>
```

```
  <Rule RuleId="PermitProvider" Effect="Permit">
   <Condition FunctionId="urn:oasis:names:tc:xacml:1.0:function:
    string-equal">
    <Apply FunctionId="urn:oasis:names:tc:xacml:1.0:function:
     string-one-and-only">
     <SubjectAttributeDesignator DataType=
      "http://www.w3.org/2001/XMLSchema#string"
      AttributeId="urn:org:chip:ping:attribute:role"/>
    </Apply>
    <AttributeValue DataType="http://www.w3.org/2001/XMLSchema#string">
     administrator</AttributeValue>
   </Condition>
  </Rule>
 </Policy>
</PolicySet>
```

B.2 Request

This is the example request sent from the client to the Indivo server. It contains attribute values for subject, resource and action. The subject has two attributes: one specifying the id of the subject, and the other what kind of role the subject holds. The resource contains one attribute that in this example is set to *dummy-resource*. The action has one attribute with attribute id set to *urn:org:chip:ping:action:create*.

```
<?xml version="1.0" encoding="UTF-8"?>
<Request xmlns="urn:oasis:names:tc:xacml:1.0:context"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
 <Subject>
  <Attribute AttributeId="urn:oasis:names:tc:xacml:1.0:subject:subject-id"
   DataType="urn:oasis:names:tc:xacml:1.0:data-type:rfc822Name">
   <AttributeValue>torstein@indivohealth.org</AttributeValue>
  </Attribute>
  <Attribute AttributeId="urn:org:chip:ping:attribute:role"
   DataType="http://www.w3.org/2001/XMLSchema#string">
   <AttributeValue>administrator</AttributeValue>
  </Attribute>
 </Subject>
 <Resource>
  <Attribute AttributeId="urn:oasis:names:tc:xacml:1.0:resource:resource-id"
 DataType="http://www.w3.org/2001/XMLSchema#string">
 <AttributeValue>dummy-resource</AttributeValue>
  </Attribute>
 </Resource>
 <Action>
```

```
  <Attribute AttributeId="urn:oasis:names:tc:xacml:1.0:action:action-id"
   DataType="http://www.w3.org/2001/XMLSchema#string">
   <AttributeValue>urn:org:chip:ping:action:create</AttributeValue>
  </Attribute>
 </Action>
</Request>
```

B.3 Response

The response below contains the decision Permit. The value of the attribute id *urn:org:chip:ping:attribute:role* and action id *urn:oasis:names:tc:xacml:1.0:action:action-id* in the request applies to the policy's rule and target.

```
<Response>
 <Result ResourceID="dummy_resource">
  <Decision>Permit</Decision>
  <Status>
   <StatusCode Value="urn:oasis:names:tc:xacml:1.0:status:ok"/>
  </Status>
 </Result>
</Response>
```

APPENDIX C

Access Control Model Example

To demonstrate our access control model, we here present a proof of concept example including a patient and his access policies. Kåre Krank is a 58 year old male suffering from a disease that causes him to see both physicians and a physiotherapist regularly. The following information shows some of the content in the PHR. In this example Ø denotes the empty set, while _denotes wild card.

C.1 Global Information

This is where information that is true for all records in the PHR are kept. This includes OPERATIONS, ROLES, INSTITUTIONS, USERS, IRA, GUA and RH. Normally IH is also a part of this, but we have omitted the institution hierarchy in order to keep the example simple.

OPERATIONS:
R - Read
RW - Read/Write
D - Deny

ROLES:
R1 - Physician
R2 - Intern
R3 - Chief Physician
R4 - Physiotherapist
R5 - Nurse
R6 - Patient
R7 - Primary Physician

INSTITUTIONS:
I1 - Hospital
I2 - Doctor office
I3 - Physiotherapy clinic

USERS:
U1 - Dr. Frisk
U2 - Dr. Sleip
U3 - Dr. Ludvigsen
U4 - Kåre Krank
U5 - Kari Hansen
U6 - Ola Jansen

IRA ⊆ ROLES x INSTITUTIONS:
(R1,I1), (R2,I1), (R3,I1), (R4,I3), (R5,I1), (R5,I2), (R1,I2), (R7,Ø)

GUA ⊆ USERS X IRA:
(U1,(R3,I1)), (U2,(R2,I1)), (U5,(R5,I1)), (U3,(R1,I2)), (U6,(R4,I3))

RH ⊆ ROLES x ROLES:
(R7,R1), (R2,R1), (R3,R2)

C.2 Record Specific Information

This is where information that is only true for one PHR record is kept. In this case it is Kåre Krank's record. The information includes RESOURCES, PERMISSIONS, GROUPS, assignments and permission assignments. Assignments include UGA, IRGA, RUA while permission assignments include GPA, UPA and IRPA.

RESOURCES:
ReA, ReB, ReC, ReD

PERMISSIONS:
(ReA,R), (ReA,RW), (ReB,R), (ReB,RW), (ReC,R), (ReC,RW), (ReD,R), (ReD,RW)

GROUPS:
G1 - Arthritis treatment

UGA ⊆ USERS x GROUPS:
(U6,G1)
Ola Jansen (U6) is part of the Arthritis treatment group (G1).

IRGA ⊆ IRA x GROUPS:
((R1,I1),G1)
All physicians at the hospital (I1) are members of the group (G1). Due to the role hierarchy this also includes all roles at I1 that inherits the role R1: Intern (R2), Chief Physician (R3).

RUA ⊆ USERS x IRA:
(U3,(R7,Ø))
The primary physician (R7) of Kåre Krank is Dr. Ludvigsen (U3).

GPA ⊆ PERMISSIONS x GROUPS:
((ReA,R),G1), ((ReB,RW),G1)
The members of the arthritis treatment group (G1) have read access to resource ReA and read/write access to resource ReB.

UPA ⊆ PERMISSIONS x USERS:
((ReC,RW), U1)
Dr. Frisk (U1) has read/write access (RW) to journal contents ReC.

IRPA ⊆ PERMISSIONS x IRA:
((ReD,RW),(R1,_))
All physicians (R1) have read/write access to resource ReD. Due to the role hierarchy this also includes all roles that inherit the R1 role: Intern (R2), Chief Physician (R3), Primary Physician (R7).

Bibliography

[1] Knut Halvor Larsen and Torstein Jensen. Access Control Model for Personal Health Record. http://www.pvv.ntnu.no/~gunnarre/studies/for/access06.pdf, Autumn 2006. Specialisation assignment TDT4700 Healthcare informatics.

[2] Connecting for Health. The personal health working group final report. Markle Foundation, July 2003.

[3] Paul C. Tang, Joan S. Ash, David W. Bates, J. Marc Overhage, and Daniel Z. Sands. Personal Health Records: Definitions, Benefits, and Strategies for Overcoming Barriers to Adoption. *J Am Med Inform Assoc*, 13(2):121–126, 2006.

[4] Terje Brasethvik and Anders Kofod-Petersen. Eigenjournal: A personal collaborative medical journal. *Norwegian Centre of Electronic Health Record*, 2006.

[5] Torbjørn Nystadnes. Legemiddelopplysninger i Samtykkebasert kjernejournal. Technical Report 29/05, KITH, 2005.

[6] Sosial og helsedirektoratet. Individuell Plan 2005 - Veileder til forskrift om individuell plan. http://www.shdir.no/vp/multimedia/archive/00005/IS-1253_5061a.pdf. Last visited: 10/05-2007.

[7] Helse og omsorgsdepartementet. Forskrift om individuell plan etter helselovgivningen og sosialtjenesteloven. http://www.regjeringen.no/upload/kilde/hod/rus/2004/0006/ddd/word/234331-vedlegg_i-17-2004.doc. Last visited: 10/05-2007.

[8] Visma Unique. Unique SamPro - En individuell plan som gir resultater. http://www.visma.no/archive/visma_no/Produktrelaterte%20downloads/Offentlig%20sektor/UniqueSamPro.pdf. Last visited: 6/11-2006.

[9] Visma Unique AS. Unique SamPro Brukerveiledning. https://www.sampro.no/TestPilot/SamProHelp/HelpSamPro.pdf. Last visited: 10/05-2007.

[10] William W. Simons, Kenneth D. Mandl, and Isaac S. Kohane. The PING Personally Controlled Electronic Medical Record System: Technical Architecture. *J Am Med Inform Assoc*, 12(1):47–54, 2005.

[11] OASIS eXtensible Access Control Markup Language (XACML) Technical Committee. eXtensible Access Control Markup Language (XACML) Version 2.0. Technical report, Organization for the Advancement of Structured Information Standards (OASIS), 1/2-2005.

[12] OASIS eXtensible Access Control Markup Language (XACML) Technical Committee. A Brief Introduction to XACML. `http://www.oasis-open.org/committees/download.php/2713/Brief_Introduction_to_XACML.html`, 14/4-2003. Last visited 11/5-2007.

[13] Phil Griffin. Introduction to XACML. `http://dev2dev.bea.com/pub/a/2004/02/xacml.html`, 19/2-2004. Last visited 11/5-2007.

[14] Children's Hospital Informatics Program. Indivo 3.0 Beta Source Code. `https://scm.chip.org/svn/repos/ping/tags/ping-3.0-beta`. Last visited: 27/05-2007.

[15] Inc. Sun Microsystems. Sun's XACML Implementation - Programmer's Guide for Version 1.2. `http://sunxacml.sourceforge.net/guide.html`, 11/7-2004. Last visited 5/6-2007.

[16] Steinar Kvale. *Det kvalitative forskningsintervju*. Gyldendal Norske Forlag AS, 1997. Norwegian transelation by Tone M. Anderssen and Johan Rygge.

[17] Michael J. Muller. *The human-computer interaction handbook: fundamentals, evolving technologies and emerging applications*, chapter 54 - Participatory design: the third space in HCI, pages 1051–1068. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, USA, 2003.

[18] Michael J. Muller. PICTIVE - an exploration in participatory design. In *CHI '91: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 225–231, New York, NY, USA, 1991. ACM Press.

[19] Morten Kyng. Making representations work. *Commun. ACM*, 38(9):46–55, 1995.

[20] David F. Ferraiolo, Ravi S. Sandhu, Serban I. Gavrila, D. Richard Kuhn, and Ramaswamy Chandramouli. Proposed NIST standard for role-based access control. *Information and System Security*, 4(3):224–274, 2001.

[21] Anna Johansson. *Narrativ teori och metod*. Studentlitteratur AB, 2005.

[22] QSR. Nvivo 7.0. `http://www.qsrinternational.com/products/productoverview/NVivo_7.htm`. Last visited: 30/04-2007.

[23] Anders Lindseth and Astrid Nordberg. A phenomenological hermeneutical method for researching lived experience. *SCANDINAVIAN JOURNAL OF CARING SCIENCES 18 (2): 145-153*, 2004.

[24] Hroar Piene, Arnt Ole Ree, and Torbjørn Nystadnes. EPJ standardisering: Dokumentasjon av forskrivning og administrering av legemidler mv. Technical Report 08/03, KITH, 2003.

[25] Gøran Sveia Kvarv and Hans Magnus Wold. Reception and Representation of Electronic Prescriptions in the Indivo Personal Health Record, Autumn 2006. Specialisation assignment TDT4700 Healthcare informatics.