# Considering Privacy and Effectiveness of Authorization Policies for Shared Electronic Health Records*

Thomas Trojer
Institute of Computer Science
University of Innsbruck
Innsbruck, Austria
thomas.trojer@uibk.ac.at

Basel Katt
Institute of Computer Science
University of Innsbruck
Innsbruck, Austria
basel.katt@uibk.ac.at

Thomas Schabetsberger
ITH-icoserve GmbH
Innsbruck, Austria
thomas.schabetsberger@ith-icoserve.com

Ruth Breu
Institute of Computer Science
University of Innsbruck
Innsbruck, Austria
ruth.breu@uibk.ac.at

Richard Mair
ITH-icoserve GmbH
Innsbruck, Austria
richard.mair@ith-icoserve.com

## ABSTRACT

A central building block of data privacy is the individual right of information self-determination, once these information identify individual persons and can therefore be considered as sensitive. Following from that when dealing with shared electronic health records (SEHR), citizens, as the identified individuals of such health records, have to be enabled to decide what medical data can be used in which way by medical professionals. In this context individual preferences of privacy have to be reflected by authorization policies enforced to control access to personal health records. We see two potential challenges, when enabling patient-controlled access control policy authoring: First, an ordinary citizen is considered a non-security expert, thus not necessarily aware of implications of her/his actions of defining access control to protect personal health data. Second, permissions to access medical data are necessary to support the daily routines of medical personnel. The better the health-care information system supports these work procedures the more effective and useful it is. There should be a balance between access restrictions through privacy settings and required access permissions in order to allow the system to be effective. In this paper we present a case study in the context of SEHR in Austria. In this scenario we identify different types of authorization policies to support individuals' privacy. Patient privacy is an important factor in access decision making, but in order to ensure the privacy – effectiveness balance, citizen-authors of policies should be informed about implications of their privacy settings on the underlying information system. To ensure this balance, policies need to be analysed. In this paper we describe a policy analysis method based on generated rules to evaluate the consequences of citizens privacy settings. Analysis results can then be used to inform and support a citizen during the policy authoring process.

*This work was partially supported by the Austrian Federal Ministry of Economy as part of the Laura-Bassi – Living Models for Open Systems – project FFG 822740/QE LaB, see http://lab.q-e.at/


*IHI'12,* January 28–30, 2012, Miami, Florida, USA.


## Categories and Subject Descriptors

D.4.6 [**Security and Protection**]: Access controls; K.4.1 [**Public Policy Issues**]: Privacy; H.1.2 [**User/Machine Systems**]: Human factors

## Keywords

Access control, Policy analysis, Electronic health record

## General Terms

Design, Human Factors, Legal Aspects, Security

## 1. INTRODUCTION

Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others [21]. Therefore data privacy defines itself as a protection mechanism to mitigate personal damage when data can be brought into the context of individual persons. A central building block of privacy is the individual right to decide which data about oneself might be collected and stored and how data is supposed to be processed [15]. Since 1983 this right, entitled *informational self-determination* is a fundamental right in German law[1] and further is a substantial part of the European Data Protection Directive 95/46/EC [5] established all over in countries of the European Union (EU) through corresponding national law.

In order to ensure privacy, national law has to be observed as well as technical considerations have to be made. Further to enforce the right on informational self-determination, management and authoring capabilities for data and policies have to be implemented and provided to the end-user. We identify two main challenges to be tackeled towards this goal. First, users are typically not security experts and the

[1]law of the German federal constitution (BVerfGE) 65, 1

actual definition or selection of privacy preferences requires the user to translate their mental conception of privacy into security configurations. Second, access control restrictions might have a negative impact on the effectiveness and the functionality of the system. The first issue has been well discussed in literature, e.g., by suggesting usable authoring tools, like the line of work done by Karat et al. [8]. However, little attention has been paid to the second problem, which is the focus of this paper.

### *Privacy-Effectiveness Balance.*

When dealing with potentially complex and inter-connected data, which is accessed by multiple stakeholders for different purposes, an individual is concerned about her/his privacy. Authoring tools should support users to encode these privacy concerns into enforceable security configurations. Besides the importance of having such tools being user friendly, they have to ensure that defined policies have no impact on system functionality and effectiveness. Further users have to be able to validate whether their conception of personal data privacy matches with the preferences enforced through corresponding policies. Being aware of the coverage of those policies a user is able to decide if adaptations to those preferences are required. On the other hand, institutional stakeholders in certain cases need to access personal information of users in order to accomplish the goals and tasks defined by their daily work. For example, in order to do a surgery, doctors needs to access the relevant medical history of the patient. Obviously enforcing a user-defined privacy policy has a potential to interfere with the tasks a stakeholder has to execute, therefore conflicting with the goals of an institution [13]. A well-balanced information flow [15, 12] is considered an important factor to support both, individual data privacy and information system effectiveness. In many scenarios the functioning of an institution is mainly based on information flow, decision making upon retrieved information and adaptations to information. On the other hand limitations to this information flow protect the individual person against exaggerated data acquisition and misuse of already collected data.

Figure 1 shows the relation between user privacy and system effectiveness with regard to access control settings *(S0...S4)*. Setting *S1* indicates that no access control restrictions are applied and resources are available to everyone. For example, a public web site that provides weather forecast services does not have to limit access because of privacy considerations. *S2*, on the other hand, indicates the setting that provide no permissions to the protected resources. These setting are useless since the functionality of the system is completely blocked by the access control mechanisms. The ultimate goal for an authoring tool for access control configurations is to ensure the best balance between system effectiveness and access control restrictions (cf. access control setting *S0*). Settings that are positioned in the area between *S0* and *S1*, like *S3*, indicate that effectiveness overweights privacy, while those that are positioned in the area between *S0* and *S2*, like *S4*, indicate more privacy restriction are applied on the cost of system effectiveness.

In order to specify access control and privacy settings, a plenty of policy standards, languages and models have been proposed, however, the concept of effectiveness is not well defined in the literature. In this paper we define the concept of effectiveness based on two factors. First, the *needs-to-know* [7] relationship between users and specific resources. For example, a pharmacist needs to know the prescription of the physician in order to deliver the right drug to the patient. Second, *personal relationship* between two users of the system. This relationship indicates a certain basis for trust, stating that one user can be allowed to access information about the other one. For example, the primary physician who has a personal relationship with a patient should be granted access to the health history of her/his patient in the context of a medical treatment.

### *Contributions.*

In this paper we tackle the important issue of finding the right balance between (*i*) access restrictions set by individuals through enforced privacy preferences and (*ii*) access needs by health-care stakeholders required to accomplish their tasks. In order to elaborate our approach, we consider a national case study. In this case study we discuss, on one hand side, all possible scenarios for stakeholders using an electronic health record of a citizen and, on the other side, the privacy and access control requirements imposed by the Austrian law. Effectiveness is defined based on two factors, namely *needs-to-know* and *personal-relationship* associations. Finally, we show how users can be supported while defining access control policies as the authoring tool reports the consequences of these policies to the effectiveness of the system.

### *Outline.*

The rest of this paper is organized as follows. In Section 2 we present our case study about the shared electronic health record (SEHR) in Austria. The policy authoring model and its specification is discussed in Section 3. In Section 4 we elaborate on the analysis of authorization policies regarding privacy and the effectiveness of a health-care information system. Therein we briefly describe our prototypical implementation and finally conclude our work in Section 6.

## 2. A CASE STUDY: SEHR IN AUSTRIA

Our use-case is related to the national e-health initiative in Austria (ELGA) which started to make progress as a governmental working group from 2006. Therein a distributed but inter-connected patient record containing all relevant health-care data (i.e. multimedia, treatment session protocols, prescriptions and medication information, discharge letters, etc.) about a patient undergoes continuing discussion and ongoing implementation.

An important goal of these development efforts regards a portal application for both, patients (i.e. all citizens) and any medical personnel to view and alter patient medical data in a location and time-independent manner. Therefore legal and general technical questions as well as questions on setting up the required infrastructure are part of that initiative. A study [6] conducted to emphasize on the technical and the legislative feasibility of ELGA shows major issues within required organizational support (i.e. executing pilot projects, adjusting or setting-up infrastructure) and protection of privacy (i.e. reaching compliance with national/European data privacy law, or regulating these data by a special law).

The use of an information system providing access to SEHR is originally motivated by lowering costs of medical treatments or medical research as well as to increase the

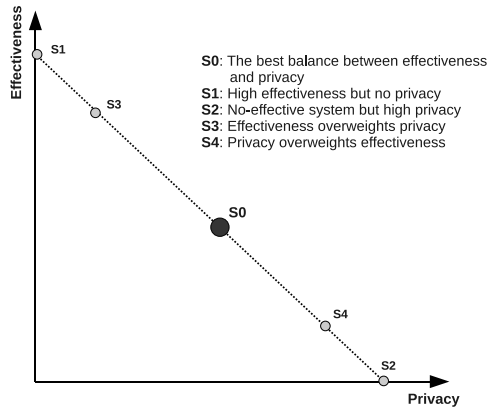


**Figure 1: The relationship between privacy and information system effectivness.**

overall effectiveness of the health-care enterprises. While effectiveness can be increased since a holistic health record of any patient can be accessed and therefore used to support a practitioner during a treatment session, privacy of patients must be protected. In the rest of this section we present use-case scenarios that identify effectiveness and privacy requirements of SEHR in Austria.

## 2.1 Use-case scenarios

In the following we list the major use-case scenarios which we encountered during discussions with our industry partner *ITH-icoserve technology for healthcare*[2] – a subsidiary company of *Siemens* and a regional hospital operator *Tiroler Landeskrankenanstalten* (TILAK). These use-case scenarios are based on settings regarding the implementation of SEHR in Austrian and according to Austrian electronic health-care and data privacy law.

### 2.1.1 Medical treatment towards an attending patient

During a treatment session, both, a patient and a medical practitioner are attending. When accessing medical records of the patient via the portal application, both parties have to state their participation in the session. Therefore two different types of smartcards together with a health-care network gateway deployed at the health-care providers side are used. The Austrian *e-Card* authenticates a patient. By providing a PIN[3] to the card reader, the patient gets authenticated and her/his actual presence is assumed. Similarly, a medical practitioner proofs her/his identity and attendance in a medical session via the *o-Card*[4].

**Access during/after treatment.** This scenario covers access for viewing or updating (i.e *data processing*) the *health record* of a *patient* during or after a *medical treatment* performed by a *practitioner*. **Permission** to access the health record is received if the type of health record relates to the assigned professional *role* of the medical practitioner performing the treatment. Further explicit permissions/**restrictions** to a medical record have to be evaluated. These may potentially be set by a citizen (**patient privacy policy**) or are stated within a **global authorization policy** (e.g., permitted access for a limited amount of time after a treatment where the health record shall stay accessible).

**Referral.** A *referral* is proposed by a *practitioner* and suggests the visit of other *medical staff* (e.g., a specialist) by a *citizen*. The practitioner referred to can be considered for **granted access** to *health record*s of the citizen prior to the proposed visit. This scenario allows for sharing medical data within limited but feasible and observable boundaries. This functionality builds on top of the strength of a SEHR, as communication of medical data does not require changes within the type of communication media anymore.

### 2.1.2 Citizen health record access

Accessing the own personal health record is possible via the portal application through a secured connection to the health-care network. Further a citizen has to authenticate her-/himself by providing the e-Card with its PIN or, if no local smartcard reader is available, with username and password credentials as a fallback option.

**Patient viewing own health record.** Here we deal with *citizens* who access and view their own personal *health record*. Access to a personal record is **restricted by** the access rules either set by the citizen (i.e. **self protection measures**) or the *practitioner* who did not release the *medical data* yet (i.e. **patient protection measures**).

The purpose of patient protections measures is explained within use-case scenario *Release of medical records*. Self-protection measures on the other side are important in situations where a citizen may not be sure if her/his medical records are read in a privacy-respecting environment. Therefore it is necessary that the identified person of a medical record is able to define the visibility of records according to e.g., **location** or **time**. Self-protection may be applied in cases where unauthorized persons are able to assert pressure and demand the disclosure of medical data or the health status of a citizen. Such unauthorized persons are e.g., a prospective employer who wants to base an engagement decision on the health status of a person. Further parents might try to observe their young adult children, which might lead to complications in their relationship in case of e.g., a recent abortion.

**Patient self care.** In this scenario a *patient* accesses her/his own *medical record* to maintain self-generated *medical data* (like blood glucose level, heart rate, blood pressure, etc.). Specific *health record type*s are available to support the self-management of medical data.

**Proposal for (non-)disclosure of health records.** Either the system, by analyzing the **patient access control policy**, or *medical staff* may propose the disclosure or even stricter protection measures of a *health record* to the identified *citizen*. This is used in case of a *referral* (see use-case scenario *Referral*) or if the system detects that either the privacy of a patient is at risk or **access restrictions** should be weakened in order to support the basic workflows of medical *personnel* (see Section 4).

[2] `http://www.ith-icoserve.com`
[3] personal identification number
[4] a smartcard for medical professionals in Austria, similar to the e-Card, but with extensional features, e.g., the retrieval of medical roles of practitioners

**Delegation of control of medical records.** A patient is allowed to **delegate** control over her/his entire *health record* at once to a related person, e.g., a relative, friend or the family practitioner. To accomplish a delegation of control a local session between both parties has to be established. This is done by providing the two identity smartcards or corresponding credentials to the system. Further delegations beyond the one(s) originally set by the citizen identified by the health record are not possible. The purpose of delegation of control is to enable the management (i.e. privacy settings, patient self care, etc.) of a health record and the participation within the health-care network even for citizens not personally able to do so. E.g., elderly or non-computer literate people may delegate the ability to maintain their medical data to their family practitioner or relatives. Further parents are declared as delegates of their non-adult children.

**Trusted medical staff.** A *citizen – medical staff* **trust relationship** shall be definable. The amount of trust of a relationship can be assigned to a practitioner by the patient via labels. Implicit to each trust label there exists a set of **access control rule**s which define how trusted medical staff can access the *health record* of a patient. Allowing for trusted medical staff is especially practical to let a citizen express her/his conception of trust roles. Common trust labels are *family practitioner* (i.e. most documents are accessible), *primary specialist* of an arbitrary medical field (i.e. documents within that field are accessible even before a medical treatment or in the long term after a treatment) or *former practitioner* (i.e. indicating a past trust relationship which prevents the practitioner from accessing documents). If no trust label is explicitly assigned a **global authorization policy** restricts access accordingly.

### 2.1.3 Medical professional access to health records

Medical professionals may access medical records without the presence of the identified patient for a dedicated purpose and under certain premises. In order to access a health record of a patient the medical practitioner has to authenticate her-/himself via the o-Card. The health-care network is only accessible from within certified locations (i.e. health-care institutions), so that remote and non-auditable access by practitioners (e.g., at home) is prevented.

**Request for patient consent.** *Medical staff* is able to request a *citizen* to be **granted** for access to her/his *health records*. Such requests only target medical records which are visibly listed but not yet viewable by the practitioner. The option for requesting a **patient consent** for access is useful if a practitioner who is currently involved in a medical treatment considers certain protected medical records as important. In a later stage this feature can be extended to ask patients to (anonymously) participate in medical studies by providing specific types of (de-identified) medical data.

**Release of medical records.** In this scenario we consider *medical staff* as the author of **non-disclosure** policies. Once such a policy protects the *health record* from being accessed by a *citizen*, only the issuing practitioner or the one a patient is *referred to* is able to release the document to the citizen. Explicit release is necessary in cases where the content of a medical record, which e.g., reflects the results of laboratory test results, contain critical parts (e.g. symptoms of a deadly disease). By placing a non-disclosure policy for a specific medical record the patient is prevented and protected from reading the health record. Finally documents can be released once further consultations have been made or further attendance of the patient at a practitioners side (e.g., as part of a referral to a specialist) has happened.

**Emergency access.** Emergency access is always performed with an access control overruling **purpose** in order to treat a patient in non-regular and critical situations. In such urgent situations access to a patient's entire *health record* shall be **granted** to *medical staff* without a proper check of access permissions. A so called breaking-glass policy [3] is therefore put in place to overrule potential access restrictions. Typically the overruling decision is coupled with the enforcement of additional **obligation**s the requester implicitly agrees to. This includes e.g., extensive logging by the system or detailed post-access reporting required to be performed by the requester. Violations to the use of emergency access or the fulfillment of obligations bound to it, have to imply legal consequences.

From these use-case scenarios we derived terms and key concepts (cf. emphasized keywords) related to the use of SEHR. Figure 2 depicts this in a compound form as a UML class diagram. Authorization policy-related concepts (cf. bold keywords) will be further described in Section 3.2.

# 3. POLICY AUTHORING FOR SEHR

As an integral part of every use-case scenario (see Section 2.1), we see the authorization mechanisms to protect personal medical data. Therefore access control enforcement as a mechanism for privacy protection supports the authorization requirements deduced from each scenario. On the other hand every use-case scenario defines requirements of having accessible medical data and allowance for flow of information. Therefore policy authoring for SEHR has to focus on both, the definition of enforceable authorization policies as well as on aspects concerning privacy and effectiveness of those policies.

## 3.1 Policy Authoring Model

The policy authoring model (see Figure 3) is a schema for describing authoring capabilities for authorization policies. Core entities for policy authoring are defined by an *Access target*, which relates an access requesting individual to the targeted resource via the desired operation to be executed on the target resource. This can be formulated by

$$Access\_target = Subject \times Resource \times Action,$$

where *Subject*, *Resource* and *Action* are the corresponding sets of available domain entities.

Based on the access target we define two models covering different usage aspects. The first aspect are enforceable authorization policies representing privacy settings. This aspect is described by the *Authorization policy model* (see Figure 3 and Section 3.2). The other aspect covers the concept of information system effectiveness influenced by authorization policies. This aspect is shown as the *Effectiveness model* (see Figure 3 and Section 4.1). Each of these aspects reasons about the access target and contributes decisions. On the one hand a privacy setting (i.e. an instance of the

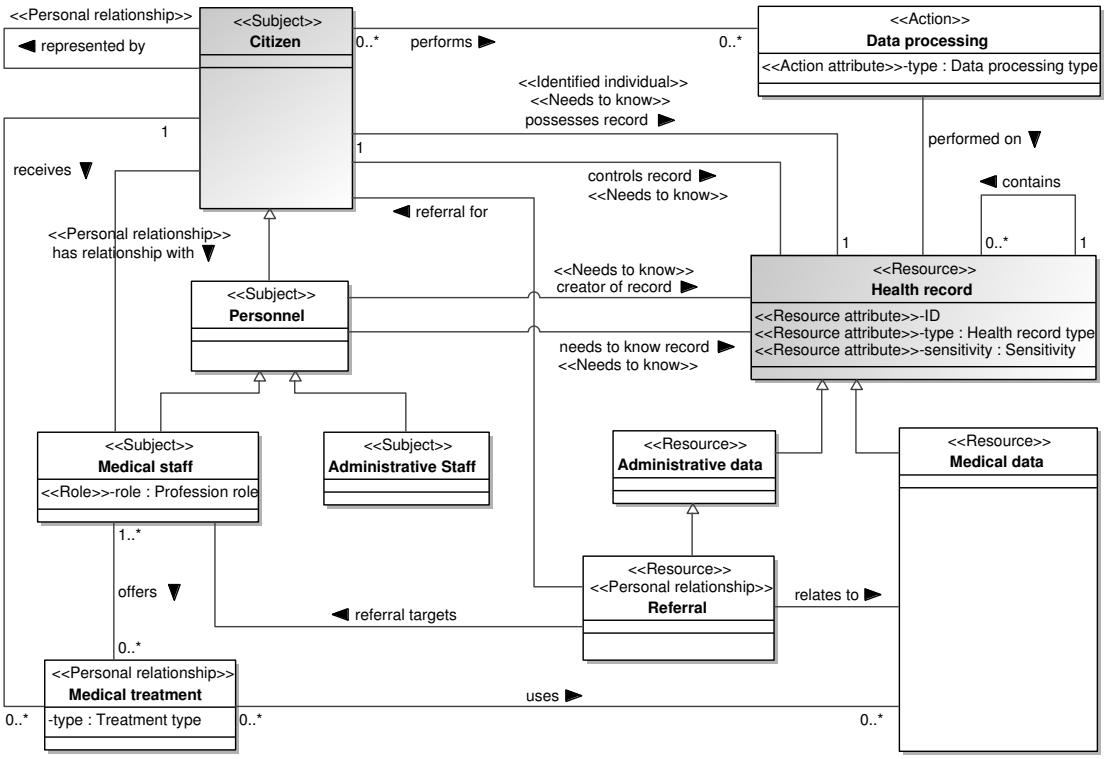


**Figure 2: Domain concept model based on entities of the health-care domain which participate in data processing of shared electronic health records.**

authorization policy model) decides whether an access request (defined by an equivalent of the access target) shall be permitted or denied. On the other hand the analysis of effectiveness (i.e. an instance of the effectiveness model) decides whether the flow of information is hindered, therefore limiting the effectiveness of an information system. The policy authoring model unifies those orthogonal aspects of policy decision-making to a common foundation for implementing authorization policy authoring applications.

The key step towards implementing policy authoring is to apply the policy authoring model to the specific domain model. In our context we derived all domain concepts from our use-case scenarios related to SEHR. The outcome of applying the authoring model to the health-care domain model is the *health-care authoring domain model*, which is indicated within Figure 2. One feature of having a health-care authoring domain model is that it represents a schema to let a policy author derive access control policies. Such policies always match the available authorization concepts and correspond to the health-care domain concepts. Further, arbitrary authorization policies can be validated against this schema to check for compliance. Another important purpose of the health-care authoring domain model with its unified integrated policy aspects is that authoring of a policy reflects all given aspects. E.g., when defining an authorization rule allowing a patient's family practitioner to access all of her/his medical data, the authoring application might as well add a note that this rule is indeed contributing to the overall information system effectiveness. On the other hand if the authoring application encounters that a medical treatment by a practitioner towards a patient took place, it might suggest or automatically add a policy rule allowing this practitioner to create new medical record or to edit existing ones. Following from that such a model allows to evaluate the appropriateness of a policy in multiple directions.

## 3.2 Authorization Policy Model

To authorize access to shared medical data we propose a combination of three types of access control policies. Therein a *global authorization policy* builds the foundation for controlling access to health records by providing default settings offering basic privacy. This global policy is derived from regulations found within national privacy and electronic health-care law. Permissions and restrictions described by this policy include

- a default time period where access to a health record after a medical treatment is allowed,
- practitioner roles according to their field of expertise and related health-care document types which they are allowed to view and alter,
- the non-disclosure of health data known to be critical to protect the corresponding patient,
- a breaking-glass type of rule to enable emergency access to a patient health record and
- a consent-based type of rule allowing access based on a provided patient consent.

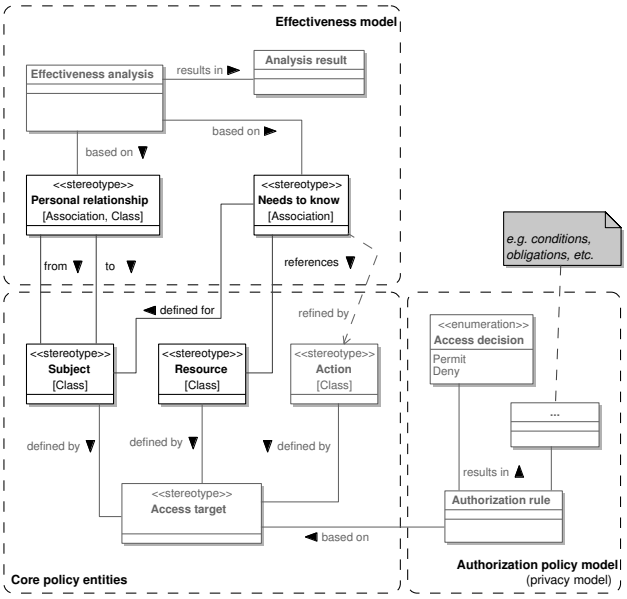


**Figure 3: Health-care authoring domain model showing core entities of an access target as well as a model for the privacy – effectiveness aspect and a model for authorization.**

Certain usage scenarios, e.g., access to data after medical treatment, are defined via the global authorization policy. Still such a policy does not necessarily match with a citizen's conception of privacy. Therefore an optional *patient privacy policy* on top of a global authorization policy reflects the needs and conception of privacy of single citizen. A patient privacy policy is defined in a way that it either extends or overwrites parts of the basic privacy principles declared by the global authorization policy. That is, the author of such a policy provides further details on her/his privacy conception or modifies parts of the globally existing privacy measures to strengthen or weaken their effects, respectively. Permissions and restrictions potentially described by this policy include

- the definition of trusted medical staff, like a family practitioner,
- the delegation of control over personal medical data to a trusted maintainer and
- citizen self-protection measures.

Separate to individual privacy access control settings on top of the global authorization policy, requests for instant access to medical data of a citizen can be made. Access is granted upon obtaining a consent of the citizen (*patient consent*). Therefore access control enforces to permit access only if an obtained patient consent can be provided. A patient consent is, different to the patients privacy policy, a temporal agreement always stating a limited amount of time and typically a limited amount of medical data (described by a set of document identifiers or via a document type) to be accessible. Permissions and restrictions covered by this policy include

- temporal access permission to a specific set of medical records approved by a corresponding citizen and

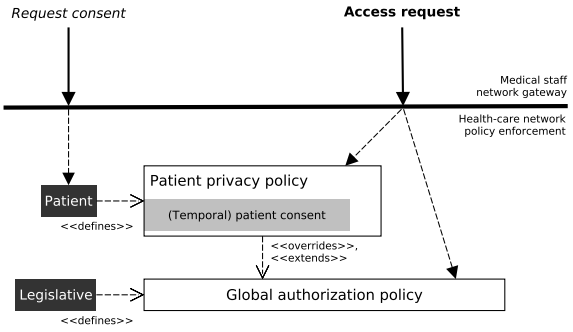


**Figure 4: Authorization policy enforcement based on a global authorization policy and a patient-defined privacy policy including temporal consents for accessing medical data.**

- emergency access (enforced by the breaking-glass type of rule) to an entire health record with the patients consent implicitly assumed.

A practical scenario of using a temporal patient consent is when practitioners want to inform themselves about a medical case before the patient is actually attending in a treatment. Further, if medical staff wants to review a health record to start new investigations for potential treatment as well as research, a request for patient consent can be issued.

The enforcement procedure, policy-defining stakeholders as well as the types of policies and their relationship is presented in Figure 4. The authorization policy model indicated in Figure 3 defines all capabilities required for protecting access to SEHR. Common entities from Role-based Access Control (RBAC) [18] or Attribute-based Access Control (ABAC) [20, 9] might be used to express authorization policies. Further, extensional attributes like *purpose* of access or purpose of collection [14, 15] of a health record can be integrated to this model. Still, the concrete specification of an authorization model is not the scope of this paper. Via UML profiling features [16] we correlate entities from the domain model to their role within policy authoring (cf. *health-care authoring domain model* described in Section 3.1). E.g., the domain model shows the entity *Citizen* functioning as the *Subject* of an access target used by an authorization rule. *Data processing* is considered an *Action*, whereas any type of *Health record* is defined to function as a potentially targeted *Resource* of an access request.

## 3.3 Authorization Policy Analysis

Authorization policy analysis is the method to check for conflicts within multiple active policies and to evaluate how a policy influences the system. In our case multiple policies (see Figure 4) are activated while accessing health-records of a single citizen. Our privacy – effectiveness balance states the definition of access control policies in accordance to the requirement of having a well-functioning and therefore effective information system. To implement this requirement we assume authorization policies to be sound in respect to their enforcement. That means, in order to enforce a policy, it has to be syntactically correct and no conflicts arise when trying to compute an access decision.

In this paper we focus on the impact of authorization poli-

cies (i.e. privacy settings) on the effectiveness of the health-care information system and vice versa. In the following section we describe the effectiveness model and show the impact on effectiveness by analysing access control policies. Analysis is done in order to achieve the goal of having privacy and effectiveness balanced.

# 4. PRIVACY VERSUS EFFECTIVENESS

Permissions to access medical data is fundamental to keep a health-care information system functioning. Functionality provided by a health-care information system like viewing medical data of patients for a treatment session or in emergency situations, as well as keeping such data up-to-date and allowing for distribution, is vital to medical personnel executing their daily routines. Based on that we define the term effectiveness by two factors: The *needs-to-know* relation between individuals and resources and the *personal-relationship* between two individuals.

## 4.1 Effectiveness Model

The *Effectiveness model* (part of Figure 3) consists of two entities to express criteria to evaluate an information system on its effectiveness and if it can be considered useful. The entity *personal-relationship* is defined between two *Subject*s and lets a policy author (or the authoring application implicitly) define an arbitrary trust relationship. Based on such a relationship privacy settings on permitted access to a corresponding access target can be justified. On the other hand the effectiveness is potentially lowered in case of explicit restrictions for trusted subjects. The *needs-to-know* relation is defined between a *Subject* and a *Resource* indicating the requirement of being able to access this resource, e.g., in order to accomplish a working task. Based on this relation it can be shown if privacy settings reflect the needs for access to resources in the context of the actual domain. By analysing both entities as part of an instance of this model, we can compute the influence and appropriateness of a patient privacy policy.

We formulate a personal-relationship by the abstract type

$T_{rel} = \mathit{Citizen} \times \mathit{Citizen}$.

As we already described in Section 3.2 regarding authorization model entities, we apply effectiveness model entities to our health-care domain model. After this application our health-care authoring domain model (see Figure 2) shows the following personal relationships based on matching the type $T_{rel}$:

$\mathit{represented\_by} = (\ \mathit{Citizen}, \mathit{Citizen}\ ) : T_{rel}$,

defines a delegation of control between the identified citizen and an arbitrary other citizen. The second-listed citizen is thereby ordered to substitute the first-listed citizen in maintaining her/his electronic health record.

$\mathit{has\_relationship\_with} = (\ \mathit{Citizen}, \mathit{Medical\_staff}\ ) : T_{rel}$,

defines an arbitrary trust relationship between a citizen and medical personnel. E.g., a family practitioner may be assigned for a specific citizen via this association.

*Medical_treatment* contains $(\ \mathit{Citizen}, \mathit{Medical\_staff}\ ) : T_{rel}$,

as a medical treatment relates a patient to an arbitrary medical practitioner. A medical treatment implicitly forms a trust relationship as defined previously. Besides the personal-relationship a medical treatment also includes associations to other entities, as indicated by the dots.

*Health_record* contains $(\ \mathit{Citizen}, \mathit{Personnel}\ ) : T_{rel}$,

as a health-record defines besides others, the entities to form a relationship between the record creating medical personnel and the citizen which is identified by this health record.

To cover the needs-to-know relationship between medical data, its association to medical personnel and also citizens which are identified by these data we define the type

$T_{know} = \mathit{Citizen} \times \mathit{Health\_record}$.

From our health-care authoring domain model we extract the following needs-to-know relationships based on the structure defined by $T_{know}$:

$\mathit{possesses\_record} = (\ \mathit{Citizen}, \mathit{Health\_record}\ ) : T_{know}$,

defines the general relationship of a citizen owning or being identified by a personal health record. As a citizen shall always be able to access her/his related data, a needs-to-know aspect is declared. The establishment of this association can be deferred by the creator of the health record as part of patient-protection measures.

$\mathit{controls\_record} = (\ \mathit{Citizen}, \mathit{Health\_record}\ ) : T_{know}$,

associates a citizen, which is delegated to maintain another citizen's health record. A citizen with such delegated control may only be allowed to take restricted actions on a health record. Still, in order to accomplish tasks of her/his maintainer role, a needs-to-know aspect is defined.

$\mathit{referral\_targets} = (\ \mathit{Medical\_staff}, \mathit{Referral}\ ) : T_{know}$,

defines a needs-to-know aspect as the targeted medical staff should be allowed to access a referral corresponding to her/him.

$\mathit{creator\_of\_record} = (\ \mathit{Personnel}, \mathit{Health\_record}\ ) : T_{know}$,

associates health-care personnel with the specific content of a health record of a citizen, which was created by her/him.

$\mathit{needs\_to\_know\_record} = (\ \mathit{Personnel}, \mathit{Health\_record}\ ) : T_{know}$,

generically describes an association between health records and personnel if the needs-to-know aspect is not implicitly justified by a connection between health-care personnel and health records (like it is e.g., in the case of the maintainer or creator of a record). This association covers requests for patient consent to disclose medical data.

To evaluate if there is a balance between privacy and effectiveness we have to investigate how authorization rules protect privacy and which value they deliver to the effectiveness of an information system. In the following section we derive analysis rules to compute the relation between privacy and effectiveness similar to what has been shown in Figure 1.

## 4.2 Derived Analysis Rules

From the authorization policy model and based on our description of entities of information system effectiveness we can derive specific analysis rules which evaluate authorization policies for their privacy – effectiveness relationship. Analysis results are then used to inform the author of an authorization policy (i.e. either a citizen or medical staff) about which aspects break the privacy – effectiveness balance. For analysing authorization policies we follow a simple

pattern, which allows us to determine a potential privacy risk or a setting hindering a potentially necessary flow of information. This pattern is based on the possibility to logically combine *Access decision* (i.e. either *Permit* or *Deny*) with a *needs-to-know* relationship (i.e. either it exists or not) and a possibly existing *personal-relationship*. Each combination is then exhaustively tested for all elements of the occurring domain entities. In the latter part of this section we focus on a meaningful subset of such combinations and label them accordingly.

Each analysis of an authorization policy together with effectiveness entities results in the notification of the corresponding policy author via the notification function,

$notify(type, weight, Access\_target)$,

where $type = \{Privacy, Effectiveness\}$, defining the type of issue encountered and $weight = \{none, inform, warn\}$, defining the required attention of the policy author towards the encountered type of issue. Formally the derived analysis rules identify the following issues:

**Strong effectiveness issue.** If a needs-to-know aspect together with a personal relationship between two stakeholders is encountered, but a restraining authorization rule (i.e. the *Access decision* is *Deny*) prohibits access to a health record, a warning regarding the information system effectiveness is created. This can be formulated as,

$(C_1, C_2) : T_{rel} \wedge (C_1, HR) : T_{know} \wedge (C_2, HR) : T_{know} \wedge$
$(AT, Deny, \ldots) \in Authorization\_rule$
$\Rightarrow notify(Effectiveness, warn, AT)$,
with $AT = (C_2, HR, \ldots) \in Access\_target$

The instance model presented in Figure 5 shows an example of the occurrence of this relationship between stakeholders and a health record.

**General needs-to-know.** If a needs-to-know aspect is present, but no personal relationship between the requesting subject and the citizen identified by the health record is given, a general statement regarding the needs-to-know aspect is created, once the subject is denied access. This can be formulated by,

$(C, HR) : T_{know} \wedge (AT, Deny, \ldots) \in Authorization\_rule$
$\Rightarrow notify(Effectiveness, none, AT)$,
with $AT = (C, HR, \ldots) \in Access\_target$

Such notification is sent e.g., if a medical practitioner requests access to a health record prior to any existing relationship with the patient (e.g., before the initial medical treatment took place). In this situation the patient gets informed that no special reason could be found, which would suggest to follow this request.

**Potential privacy issue.** A potential privacy issue is encountered if a personal relationship aspect is given together with a permission to access a health record where no needs-to-know aspect is present.

$(C_1, C_2) : T_{rel} \wedge \neg(C_2, HR) : T_{know} \wedge (AT, Permit, \ldots) \in$
$Authorization\_rule \Rightarrow notify(Privacy, none, AT)$,
with $AT = (C_2, HR, \ldots) \in Access\_target$

**Weak privacy issue.** A weak privacy issue is encountered if there is a need to know about a health record and a permission is stated, but no direct personal relationship is present. This is basically the opposite of the *general needs-to-know* rule.

$(C_1, HR) : T_{know} \wedge (C_2, HR) : T_{know} \wedge \neg(C_1, C_2) : T_{rel} \wedge$
$(AT, Permit, \ldots) \in Authorization\_rule$
$\Rightarrow notify(Privacy, inform, AT)$,
with $AT = (C_2, HR, \ldots) \in Access\_target$

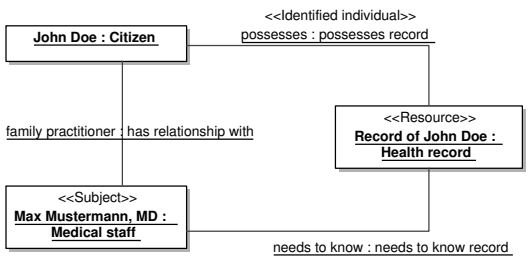


**Figure 5: An example health-care effectiveness instance model depicting two related stakeholder having a needs-to-know aspect on a health record.**

**Strong privacy issue.** If a permission is stated which allows medical staff to access arbitrary data of a patient they are not related to, the author of the authorization rule is warned.

$(C_1, HR) : T_{know} \wedge \neg(C_2, HR) : T_{know} \wedge \neg(C_1, C_2) : T_{rel} \wedge$
$(AT, Permit, \ldots) \in Authorization\_rule$
$\Rightarrow notify(Privacy, warn, AT)$
with $AT = (C_2, HR, \ldots) \in Access\_target$

## 4.3 Prototypical Implementation

We have implemented a stand-alone web-based application for access control policy authoring in the context of our health-care authoring domain model. Further a control panel for simulation purposes has been created which lets us define person-relationship aspects, e.g., via creating a referral or announcing a medical treatment which (virtually) had taken place. Integrated within a real health-care information system such data would be provided automatically, in our scenario we had to establish functionality to simulate routines of the health-care domain. As we have elaborated in [19] we automatically derived a knowledge base and corresponding interfaces from our health-care domain authorization model. Via these interfaces the authoring application alters the knowledge base with facts about access control policies, needs-to-know aspects and personal relationships between stakeholders. By using a logic programming language[5] we are able to reason about our knowledge base in a way that it represents our derived analysis rules (see Section 4.2). After every committed change to the underlying knowledge base, feedback about the analysis result is provided to the user of the application. The provided feedback is designed in a way to attract the user's attention and presents privacy or information system effectiveness issues in a human-readable way.

The screenshot depicted in Figure 6 shows effectiveness warnings to the policy author, after she/he added an explicit restriction (i.e. an authorization rule). Several facts

[5]we use SWI-Prolog, see `http://www.swi-prolog.org`

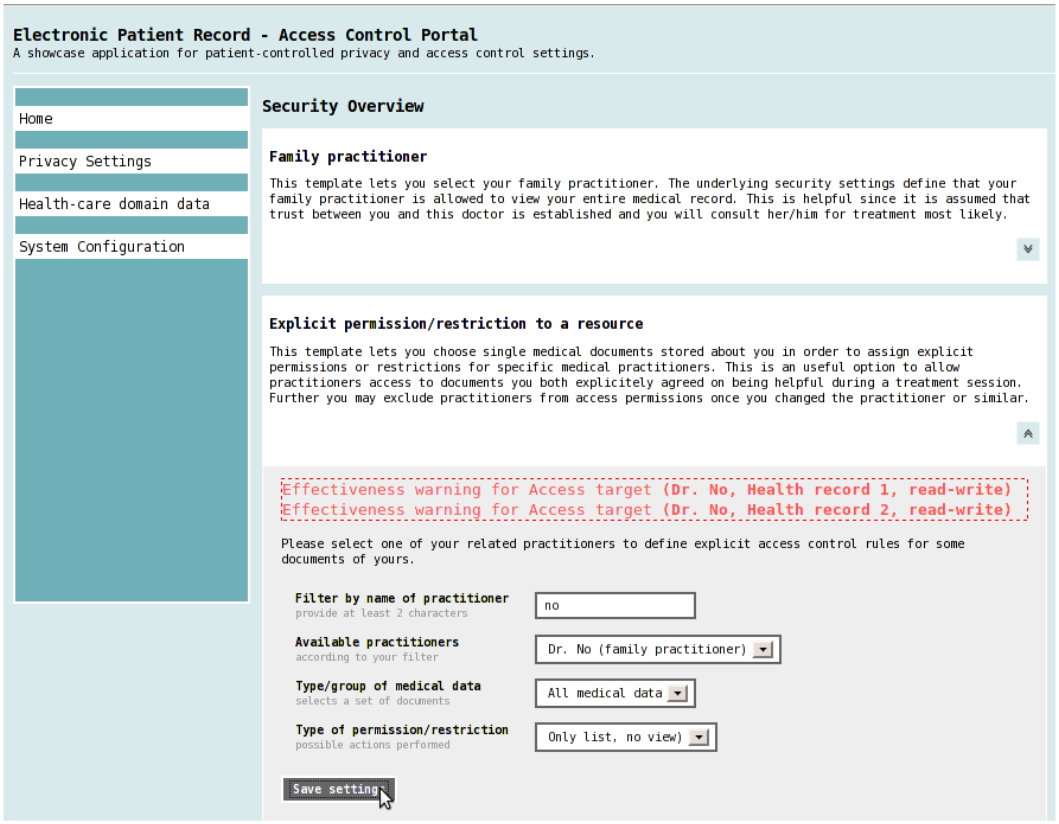


**Figure 6: Screenshot of our authoring application prototype. The content of the red dashed box indicates warnings regarding effectiveness as the author restricted all access to medical data for her/his family practitioner.**

and rules within the knowledge base lead to this notification result. A patient, namely *Patient 1* is used by the prototype for the current authoring session. This patient holds two health records, simply titled *Health record 1* and *Health record 2*. *Dr. No* is a medical practitioner and was formerly assigned as the family practitioner of this patient. By performing this assignment Dr. No was given a personal relationship to Patient 1 as well as needs-to-know aspects set on all health records of Patient 1. These multiple needs-to-know aspects arise because of the type of personal relationship that has been set. In our case setting medical personnel as the family practitioner results in a needs-to-know aspect for the entire health record of the patient. Therefore we can find, besides others, the following facts and rules (corresponding to the model in Figure 2 and stated as PROLOG code) in the backing knowledge base:

```
medical_staff('Dr. No').
citizen('Patient 1').
has_relationship_with('Patient 1', 'Dr. No').

health_record('Health record 1').
possesses_record('Patient 1', 'Health record 1').
needs_to_know_record('Dr. No', 'Health record 1').

health_record('Health record 2').
possesses_record('Patient 1', 'Health record 2').
needs_to_know_record('Dr. No', 'Health record 2').

deny(target('Dr. No', 'read-write', 'Health record 1')).
deny(target('Dr. No', 'read-write', 'Health record 2')).

notify('Effectiveness', 'warn', target(S1,A,R)) :-
    possesses_record(S2,R), needs_to_know_record(S1,R),
    has_relationship_with(S2,S1),
    deny(target(S1,'read-write',R).
```

## 5. RELATED WORK

In [7], a general proposal for implementing access control in distributed electronic health-care networks, the authors highlight the need of patient privacy policies in order to lawfully process and communicate medical information, based on a patients independent and informed decision to do so. Further the needs-to-know principle is described in this work to allow the definition of access requirements to support typical usage scenarios within the health-care domain. In our work we designed a method to actually evaluate this principle together with a patient's need for data privacy. In [17] requirements and an initial model for patient-controlled access control using RBAC is presented. Additionally the work in [2] discusses access control for medical records maintained by electronic information systems. The authors proposed, similar to our work, several models which define concepts of security related to the health-care domain.

Fundamental work on policy analysis is done in [11, 4,

13]. In [13] the authors discuss the need for policy conflict detection among authorization and imperative policies, therefore showing the importance of domain-related information to support the functioning of an information system. The authors in [1] propose logic programs to reason about access control models, in their context mandatory and discretionary models. In our work we use a similar approach to perform reasoning of authorization policies. The distinction to our work is the level of application for reasoning. While the work in [1] suggests to evaluate the model itself, we derive higher-level aspects from the model to be evaluated. Another framework for logics-based analysis of policies can be found in [10].

# 6. CONCLUSION AND FUTURE WORK

In this paper we presented a case study about the Austrian initiative to establish the electronic health record. This initiative plans to implement an important privacy concept, namely information self-determination, which places the citizen in a central position for setting privacy preferences. We modeled important entities based on usage-scenarios of SEHR and added concepts for providing authorization mechanisms. Since a citizen is not considered a security expert, nor an expert on required information flows within the health-care domain, the authoring of privacy settings is a critical task. With our approach of balancing privacy and information system effectiveness (i.e. allowing the flow of information) we contribute a step towards citizen-centric control of personal electronic health records.

Future work in this project will allow us to integrate the access control policy authoring tool into the health-care information portal developed by our industry partner. We will extend the notification of analysis results to become more usable by means of readability and by means of automatically derived problem resolution features a citizen can use. Finally we will investigate conditions further restricting authorization rules. By omitting conditions (as we did in this work) we potentially overestimated the need for notifying a citizen-author of access control rules. Still the consequences of these overestimations are not severe as conditions only limit the applicability of an access rule, but do not change its defined access decision.

# 7. REFERENCES


[1] E. Bertino, B. Catania, E. Ferrari, and P. Perlasca. A logical framework for reasoning about access control models. In *Proceedings of the sixth ACM symposium on Access control models and technologies*, SACMAT '01, 2001.

[2] B. Blobel. Authorisation and access control for electronic health record systems. *International Journal of Medical Informatics*, 73(3), 2004.

[3] A. D. Brucker and H. Petritsch. Extending access control models with break-glass. In *Proceeding of the 14th ACM symposium on Access control models and technologies*, SACMAT '09, 2009.

[4] R. Chadha. A Cautionary Note About Policy Conflict Resolution. *MILCOM*, 0, 2006.

[5] European Commision. *Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data*, 1995.

[6] Integrating the Healthcare Enterprise (IHE). *Feasibility Study for implementing the electronic health record (ELGA) in the Austrian health system*, 2006.

[7] Integrating the Healthcare Enterprise (IHE). *IT Infrastructure Access Control (White Paper)*, September 2009.

[8] C. Karat, J. Karat, C. Brodie, and J. Feng. Evaluating interfaces for privacy policy rule authoring. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, New York, NY, USA, 2006. ACM.

[9] D. R. Kuhn, E. J. Coyne, and T. R. Weil. Adding Attributes to Role-Based Access Control. *IEEE Computer*, 43(6), june 2010.

[10] M. LeMay, O. Fatemieh, and C. A. Gunter. PolicyMorph: interactive policy transformations for a logical attribute-based access control framework. In *Proceedings of the 12th ACM symposium on Access control models and technologies*, SACMAT '07, 2007.

[11] E. Lupu and M. Sloman. Conflicts in Policy-based Distributed Systems Management. *IEEE Transactions on Software Engineering*, 25, 1999.

[12] F. Massacci, J. Mylopoulos, and N. Zannone. A Privacy Model to Support Minimal Disclosure in Virtual Organizations. In *In Proceedings of the W3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement*, 2006.

[13] J. D. Moffett and M. S. Sloman. Policy conflict analysis in distributed system management, 1993.

[14] Q. Ni, A. Trombetta, E. Bertino, and J. Lobo. Privacy-aware role based access control. In *Proceedings of the 12th ACM symposium on Access control models and technologies*, SACMAT '07. ACM, 2007.

[15] OECD. *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, 1980.

[16] OMG. *UML 2.3, Infrastructure Specification*, 2010.

[17] L. Røstad. An Initial Model and a Discussion of Access Control in Patient Controlled Health Records. In *Third International Conference on Availability, Reliability and Security: ARES 2008*, Washington, DC, USA, 2008. IEEE Computer Society.

[18] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman. Role-based access control models. *Computer*, 29, 1996.

[19] T. Trojer, B. Katt, F. Wozak, and T. Schabetsberger. An Authoring Framework for Security Policies: A Use-case within the Healthcare Domain. In *eHealth 2010*, 2010.

[20] L. Wang, D. Wijesekera, and S. Jajodia. A logic-based framework for attribute based access control. In *In 2nd ACM Workshop on Formal Methods in Security Engineering (FMSE 2004*. ACM Press, 2004.

[21] A. Westin. *Privacy and Freedom*. Atheneum, New York, USA, 1967.