INTERNATIONAL JOURNAL OF MEDICAL INFORMATICS 78 (2009) 141–160

journal homepage: www.intl.elsevierhealth.com/journals/ijmi

Review

Inter-organizational future proof EHR systems A review of the security and privacy related issues

Helma van der Linden[a,*], *Dipak Kalra*[b], *Arie Hasman*[c], *Jan Talmon*[a]

[a] *School for Public Health and Primary Care: Caphri, University Maastricht, Maastricht, The Netherlands*
[b] *CHIME, University College London, London, United Kingdom*
[c] *Department of Medical Informatics, University of Amsterdam, Amsterdam, The Netherlands*

ARTICLE INFO

Article history:
Received 23 November 2007
Received in revised form
30 June 2008
Accepted 30 June 2008

Keywords:
Computerized Medical Records Systems
Data security
Access policy
Standards
Networked care

ABSTRACT

Objectives: Identification and analysis of privacy and security related issues that occur when health information is exchanged between health care organizations.
Methods: Based on a generic scenario questions were formulated to reveal the occurring issues. Possible answers were verified in literature.
Results: Ensuring secure health information exchange across organizations requires a standardization of security measures that goes beyond organizational boundaries, such as global definitions of professional roles, global standards for patient consent and semantic interoperable audit logs.
Conclusion: As to be able to fully address the privacy and security issues in interoperable EHRs and the long-life virtual EHR it is necessary to realize a paradigm shift from storing all incoming information in a local system to retrieving information from external systems whenever that information is deemed necessary for the care of the patient.

Contents

* *Corresponding author at*: Medical Informatics, University Maastricht, PO Box 616, 6200 MD Maastricht, The Netherlands. Tel.: +31 433882235.
E-mail address: h.vanderlinden@mi.unimaas.nl (H. van der Linden).
1386-5056/$ – see front matter © 2008 Elsevier Ireland Ltd. All rights reserved.
doi:10.1016/j.ijmedinf.2008.06.013

1. Introduction

Most electronic medical record (EMR) systems that are currently in use are built and implemented with only local usage in mind. Communication between healthcare workers is translated to communication between systems and implemented as a one-to-one exchange of messages where the initiating party a priori knows the party to be queried (e.g. the physician entering a lab order in the EMR system which passes the order onto a known laboratory system and the lab returning the results to the EMR system). Such an exchange is comparable to a telephone call or its Internet-equivalent: the email conversation. Tailoring the software to the specific exchange solves current interoperability issues such as differences in data structures and ambiguous interpretation due to implied metadata (e.g. absent measurement units or usage of terminology codes without specified coding schemes).

Problems arise when a new party requests access to the EMR information, the total number of users increases, the need for record sharing increases, and also when the clinical structure, organization and content need to evolve. These challenges require a generic interface that can comply with these changes, and mechanisms to find the location of the requested data.

Since the early years of this century the view has developed that high quality health care can be delivered only when all pertinent data on the health of a patient is available to the clinician. This changing point of view brings forth the notion of the (virtual) electronic health record (EHR) and requires ubiquitous communication between systems. Architectures have been designed to incorporate first generation interoperability standards, mainly in the 1990s, such as by the Synapses and Synex projects [1–3], but usually with the focus on a single organization or a single system available for multiple parties (e.g. a regional system). The HARP project [4] was the first to develop a more comprehensive componentized architecture that included a secure approach to clinical data sharing in a distributed environment. New-generation standards are better able to support elaborate communication between EMR systems. The most important are HL7 Version 3 [5] and CEN/TC215 13606 [6–10], which are now the subject of new implementation projects and technical evaluations.

Another project started by the *open*EHR Foundation is to develop an open, interoperable health-computing platform, of which a major component is clinically effective and interoperable EHR systems [11].

"The virtual EHR" is a term commonly used to denote the logical integration of distributed systems containing electronic medical or health records, irrespective of how this is achieved physically. Another commonly used term is "lifelong patient record", focusing on the availability of data and its integration over time. A lifelong virtual EHR (lvEHR) is currently seen as the best solution to meet the increasing demands for shared information described before in [12–14].

Consider the situation where an lvEHR integrates EMR systems from various healthcare organizations through the implementation of the new communication standards. Communication between these systems occurs based on information needs of the healthcare worker. The actual location of the information becomes transparent and of less importance.

Security and privacy related issues are more important in such an environment than in the current systems. For example, in the current situation access rights are defined locally, based on formal or less formal rules of the house. When dealing with access from outside the situation may ask for different requirements. In the current situation, outsiders can only have access to patient data through human intermediation, often the treating physician; she may act as a filter on what is communicated taking the patient's wishes into account. In a fully digital communication patient's wishes regarding disclosure of information have to be respected as well.

It is our objective to analyze these issues in the context of the lvEHR. We will also identify issues that require further study and regulations before safe and trustworthy lvEHR systems can become reality.

First we present the methods used, next we provide some definitions of concepts that are relevant in this article and the environment we assume in this article. Then we describe the results of our analysis, we discuss the results and present our conclusions.

2. Methods

Rather than approaching the problem of discovering and analyzing inter-organizational issues from a theoretical point of view, we have grounded the issues in a realistic clinical setting. We have used an imaginary, yet realistic scenario that focuses on shared care (see Appendix A). The scenario was divided into steps. Each step is an action that involves information exchange.

In several brainstorm sessions, H.v.d.L. and JT discussed each step from an implementation point of view (focusing on a global, rather than a more technical level). During these sessions questions were formulated that would reveal the issues without giving a solution. Defining solutions for the issues raised in earlier questions could influence later questions, a situation we tried to avoid.

The questions were grouped together in themes and answers were discussed in more detail. The arguments were supported by references to relevant literature.

We verified that our question set was complete enough by mapping the themes to topics mentioned in various standards on security and privacy in EHRs.

3. Definitions

Although there is no consensus on the exact definition of the EHR, the ISO/TS 18308 [15] standard does give a definition of the primary *purpose* of the EHR:

> The primary purpose of the EHR is to provide a documented record of care which supports present and future care by the same or other clinicians. This documentation provides a means of communication among clinicians contributing to the patient's care.

A formal definition of the scope and purpose of the Integrated Care EHR has more recently been published as ISO TR 20514 [16].

> "a repository of information regarding the health status of a subject of care in computer processable form, stored and transmitted securely, and accessible by multiple authorised users. It has a standardised or commonly agreed logical information model which is independent of EHR systems. Its primary purpose is the support of continuing, efficient and quality integrated health care and it contains information which is retrospective, concurrent and prospective."

3.1. *Definition of the EHR information*

From the ISO definitions above we can infer that the total amount of health-related information about a patient that is stored in various systems, constitutes his (lifelong) (virtual) electronic health record. The healthcare worker who is involved in the care of a specific person will be referred to as the user in this article. This user needs access to the information in the EHR to provide adequate care to the patient. This set of data is restricted by three partially overlapping divisions:

- Origin of the information: information is stored either in a local system or the information resides elsewhere (i.e. an external system). This is referred to as *local* and *external* information in this paper, respectively.
- Access: several standards [9,15] define access restrictions to protect the privacy of the patient. This is referred to as *allowed* and *unallowed* information.
- Necessity or relevancy: access to information is only allowed when it is relevant to the care process [17], again to preserve patient privacy. This is also known as the "need to know" principle.

From the viewpoint of the user, he is confronted with these three divisions: local information is readily available, while external information has to be retrieved from systems where he does not necessarily have direct access rights; he should be able to view information that is relevant to the current problem of the patient, but some of this information might be hidden due to access restrictions.

A true transparent lvEHR system would hide these divisions from the user. This allows the user to focus on his task while the system handles the burden of locating and retrieving information.

3.2. *Definition of the environment*

The environment described in the introduction resembles the configuration in Fig. 1, in which multiple organizations are connected through a central service. Each organization has a

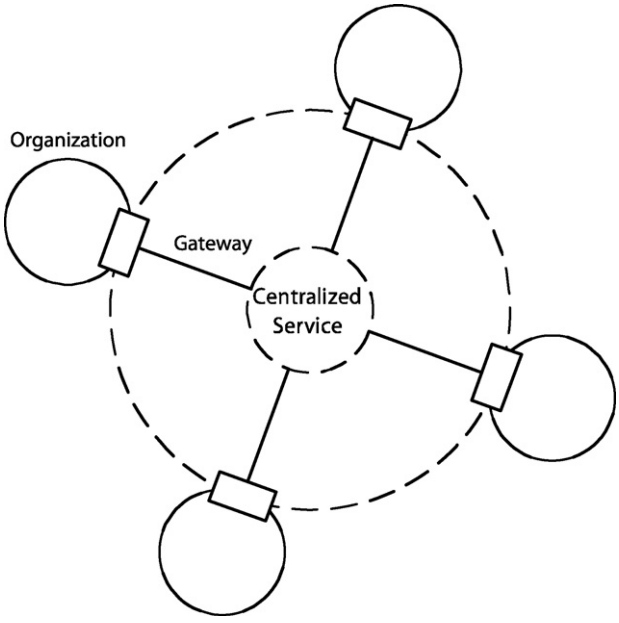

Fig. 1 – Cross-organizational environment.

gateway that hides the internal architecture from the outside world and vice versa.

The main purpose of the central service is to provide (a pointer to) the requested information to the (gateway of the) requesting EHR system(s).

This situation is comparable to an environment of distributed databases, with the difference that in a distributed database environment, the information is accessed on the database level (i.e. database schema and location are known beforehand), whereas in the environment described here the information is accessed on the application interface level (i.e. the application manages the actual information retrieval from the underlying database).

We have chosen this architecture as a basis for our discussion because it represents current best practice in a networked environment: a generic or standardized interface reduces the number of variations in implementations, while a central service reduces the number of possible connections to be made. Note that we have made no assumption about the primary location of the retrievable information, whether it remains in the respective systems or in the centralized service.

3.3. *EHR requirements*

Various authors have described generic requirements for EHR systems. We discuss only those related to security and privacy issues, with their definitions as stated in ISO/TS 18308 [15].

- *Security*: There is consensus on the security requirements for EHR systems and for communication between systems [15]. We list the definitions of the most relevant security requirements:
  - *Authentication*: verifying the claimed identity of an entity (either a person or a system).
  - *Authorization*: granting rights, which include the granting of access based on access rights, either in a personal capacity or conferred on the basis of a role held by the user.
  - *Integrity (of data)*: preserving the accuracy and consistency of data regardless of changes made.[1]
  - *Non-repudiation*: allowing any actor to obtain proof, which cannot be forged, that confirms the integrity and origin of a data item.[2]
  - *Confidentiality*: the property of data that indicates the extent to which these data have not been made available or disclosed to unauthorized individuals, processes, or other entities.
  - *Consent*: obtaining, recording and tracking of informed consent of patients for the purpose of creating and allowing access to their health information for specified purposes during specified time frames.[3]
- *Semantic interoperability*: enabling the ability to share data between systems that can be understood at the level of formally defined domain concepts to support automatic processing of data at the receiving system.
- *Author responsibility*: ensuring that each record contribution can be attributed to an identified author.
- *Audit trail or audit log*: recording activities of information system users in chronological order, which enables prior states of the information to be faithfully reconstructed. It should contain information about access to and modifications of data as well as the nature of each access and/or modification. It should also be able to support accountability for each interaction with the system by an actor (for example, logging each recorded step or task in the clinical or operational process).
- *Version management* or *version control*: supporting versioning at the granularity at which information is attested and supporting measures to discern modification or updating of the record.
- *Patient access*: allowing the patient access to all his EHR information subject to jurisdictional constraints.
- Archiving and data retention:
  - *Data retention*: providing the functionality to store patient information for at least the duration specified in legal data retention policies. This also includes the capability of preventing deletion during the retention time and of insisting on deletion of information after that period [18].
  - *Archiving*: moving EHR information to off-line storage in a way that ensures the possibility of restoring them to on-line storage when needed without the loss of meaning.

These issues can often be managed quite simply when dealing with data residing in systems in a single organization, but for the lvEHR, these issues may have other dimensions and solutions may be less trivial. Perhaps the most challenging ones relate to achieving consensus on what, how and when different requirements should be implemented.

4. Results

We started with a narrative scenario (see Appendix A for the full text) that describes a sequence of events in the common environment of paper-based forms and records dealt with by means of personal contacts through visits, telephone/fax and (e)mail. Each identified event or step was given an identifier to which we will refer in the remainder of this article. In the following we will refer to the EMR system of a specific health care provider by his acronym, followed by EMR (e.g. the EMR of the dermatologist is referred to as DERM-EMR). Fig. 2 shows a sequence diagram for this scenario.

[1] Although this is a slightly paraphrased definition of the ISO standard, integrity should not be mixed with accuracy. While the latter deals with how well a data item represents reality, integrity deals with maintaining the proper representation of the recorded data item.

[2] Although this requirement is often interpreted as being able to get proof of integrity and origin of the data that is communicated, full trustworthiness is only obtained when such proof can also be obtained on requests and receipt confirmations.

[3] The ISO standard does not define the concept of consent. We inferred this definition from the descriptions of managing consent.

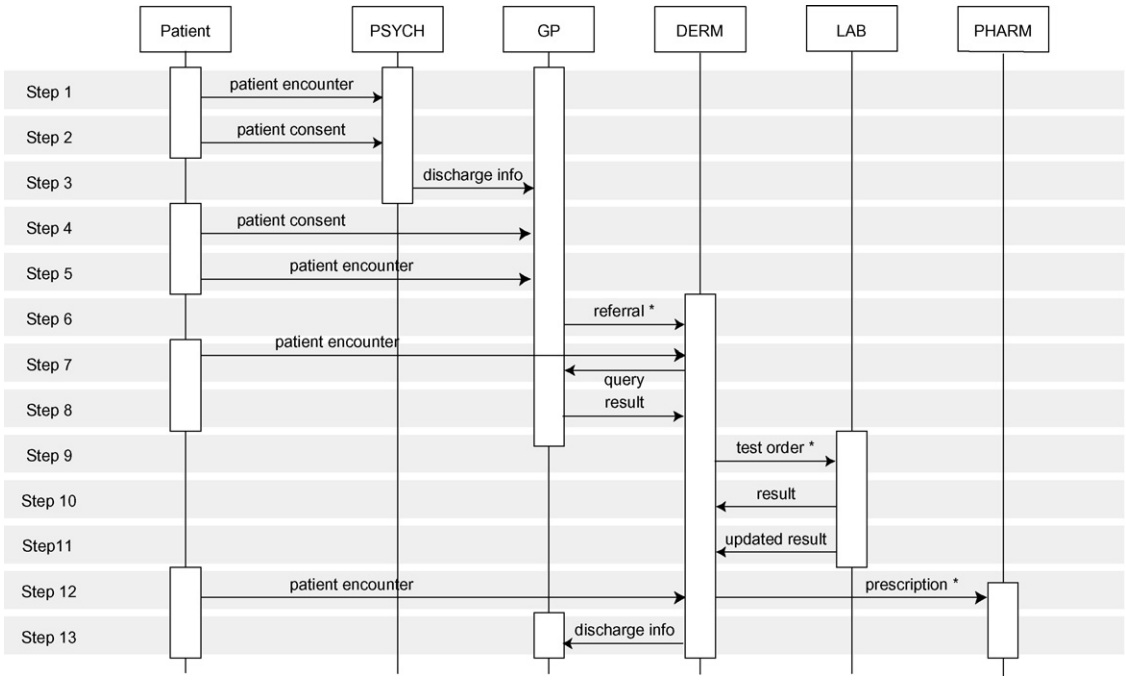

Fig. 2 – Sequence diagram of scenario. In this diagram the various columns represent the respective actors in the scenario (PSYCH = psychiatrist, GP = General Practitioner, DERM = Dermatologist, LAB = Laboratory, PHARM = Pharmacy). '*' refers to implicit patient consent.

4.1. *Issues arising from the scenario*

Analysis of the scenario resulted in a set of questions. For reference, the entire list of questions is included in Appendix B. The questions were then grouped and rephrased to reveal the underlying issues:

- *Authorized access*: how to implement authorization across organizational boundaries? (Questions 1, 2, 6, 7, 8 and 11).
- *Confidentiality*: how can it be determined that no confidentiality breach has occurred when a copy of the information resides in another system? (Question 10).
- *Patient consent*: should patient consent be given implicitly or explicitly? Can patient restrictions be overridden by emergency procedures? Can patient restrictions be used to implement "deletion" of information? Can existing patient consent be extended to a new receiving party? Can a patient grant additional ad hoc access permissions to their EHR to meet unanticipated healthcare situations? (Questions 3, 4, 12, 13 and 14).
- *Relevancy*: how to define what information is relevant and when? (Questions 3, 9, 16, 17 and 24).
- *Ownership of information*: who is the owner of information and what are the implications for exchanging information across organizational boundaries? (Questions 19 and 23).
- *Infrastructure*: what are the implications of sending correction notifications on data that is passed on to third parties? (Questions 15, 18, 20, 22 and 25).
- *Audit log*: what are the function and content of an audit log and the consequences for implementation in EHR systems? (Question 21).
- *Archiving*: what are the implications of data retention policies on the EHR content? (Question 26).

4.2. *Discussions of the issues*

4.2.1. *Authorized access*

Related questions:

1. How should a patient be identified reliably across organizations?
2. How should health professionals be identified reliably across organizations? How should organizations be reliably identified?
6. How should the PSYCH-EMR system define authorization of the GP to access information in the system?
7. Should all systems have authorization information for all possible users (i.e. persons requesting information)?
8. Should all systems provide similar access for all possible users (i.e. a GP has access to the same kind of information in all systems)?
11. In case the information from the PSYCH-EMR is stored in the GP-EMR system and matches a future query from an external system (e.g. the query from the DERM-EMR), should the information be passed on if patient consent permits or should external information always be excluded from a result set?

EMR systems contain sensitive information. It is a common requirement that access to such systems needs to be restricted to authorized users.

Reliable authorized access to patient information has three components: reliable patient identification, proper authentication of the health care provider and correct authorization of that provider.

4.2.1.1. *Identification.* Authorized access starts with a correct identification of both patient and healthcare provider. Within a single organization, it is sufficient to assign a unique code to a patient. Beyond the organizational borders, reliable patient identification across systems can be done in two ways:

- The identifications in various systems are retained and a Master Patient Index (MPI), a central service, is used to resolve collation [19].
- A national health ID (NHID) is issued by a governmental institution. In general these NHIDs are issued at birth, but can also be issued when someone starts to use the services of a health care system, for example in case of immigration. All identifications in the various systems should be replaced by or mapped to the national health ID.

The healthcare provider needs to be identified as well. Recently the Dutch government has issued a national health provider ID (UZI) [20] which can be issued to health care providers as well as to health care organizations. This ID is intended for unique identification of health providers in cross-organizational information exchange.

The main problem of the MPI is the correlation: the reliable mapping of various patient identifications to a single physical person. Studies have shown that it is not possible to fully automate this process [21].

The NHID is often described as a solution to the correlation problem, since it would require a one-time investment at the organizational level. It would also render an MPI superfluous. Many, mainly European countries, therefore regard the NHID as the more favorable approach.

Three issues arise in an environment with an NHID:

- The implications of identity theft and privacy breaches are larger.
- When patients cross national borders for health care, they will either receive a second (country-specific) ID or the local systems have to be capable of using the ID provided by the home country of the patient. Since the latter case is less realistic than the former on a pan-European scale in the present climate (due to difficulties in achieving consumer acceptance of such large scale identity sharing), an MPI to connect foreign and home IDs is still necessary.
- Patients without a national ID still need treatment and therefore a temporary ID, which needs to be correlated or replaced later with his or her national ID.

These issues would warrant the choice for an MPI, even in an environment where an NHID exists. The NHID can be seen as a solution to reduce the correlation problem, not a solution to bypass the MPI.

National IDs for healthcare professionals enable the reliable handling of situations of information exchange and are necessary to make it possible for patients to grant or deny access to their information for a specific healthcare professional (see Section 4.2.3). A unique ID at organization level would be insufficient to handle this situation.

These IDs should already be issued to students attending medical school; from the moment there is a need to access patient information, which might be much earlier than required by current policies.

Health care organizations should also be uniquely identified. This can be used to mark the original creation location of data. Care should be taken that the organizational identification is not used to construct a unique ID for the healthcare professionals, since it would falsely imply a static relationship between the organization and the professional.

Finally, unique identification should be extended to all objects, ranging from information items and system components via documents to devices and systems that are used in obtaining and maintaining health care information. This is a standard IT approach as illustrated by the IMEI number (International Mobile Equipment Identity) in cell phones and the DOI (document object identifier) used to identify scientific publications. Relevant standards of the GS1 organization [22] could be used.

4.2.1.2. *Authentication.* A centralized authentication service is common practice in an organization, since it provides various levels of efficiency. Extrapolation to a distributed environment however, raised the question of how to authenticate external access. This can be done in two ways:

- The user is known beforehand, i.e. his identification credentials are registered in the system he wants to access. This implies a separate procedure to register the credentials.
- The system relies on credentials that are issued by other systems, either from a different organization or from a general governmental body. The Dutch BIG registry, which registers all health professionals [23], is an example of the latter.

Central registries, such as the BIG registry, can be used to grant or revoke credentials for a specific professional based on his professional behavior. Legal procedures for updating and querying the status of the credentials should be created.

4.2.1.3. *Authorization.* After being granted access to the system, various levels of authorization are still necessary. Authorization can be defined as access to data or as access to functionality of the system. In the first definition, access levels vary from access to the entire health record of a patient to fine-grained access definitions at the level of medical concepts. In the latter case, users are authorized to perform specific functions (e.g. order medication, Admission Discharge Transfer). Authorization to functionality implies authorization for access to data that is necessary to use the functionality, which in turn requires semantic interoperability. For the discussion at hand, the exact distinction is irrelevant.

Each user of a system may have several roles, usually a combination of static or structural roles, based on their

position in the organization, and functional roles, based on their participation in the care process for a specific patient. Operations on resources (whether tasks or information objects) are defined as permissions. The basic relationship between roles and operations underpins Role-Based Access Control (RBAC). The set of operations that a certain role can perform on specific resources can be defined and extended through a policy [24–26]. A policy can also include constraints, both static and temporal, to avoid conflicting accumulation of permissions and to comply with regulations of various kinds. The use of policies extends the capability of RBAC to handle richer specifications of privileges. ISO TS 22600 (Privilege Management and Access Control) defines a formal architectural approach to represent such policies, and for the services that need to be implemented to support their use within a distributed computing environment [27,28].

Several sources use the term profile, but fail to sufficiently define it [9,24,29]. We define a *profile* as the set of constraints on the permissions assigned to a user, usually represented through their (structural and functional) roles and as a corresponding set of policies.

Several issues can be discerned:

- The higher the number of resources (whether tasks or information), the more complex the RBAC management will become. Evered and Bögeholz [30] have shown that RBAC implementation for a simple EHR system (nursing home with some 30 residents, 5 different structural roles and 4 sets of information) is already complex. Lovis et al. describe an implementation of an RBAC model in the Geneva University Hospitals [31]. Although multiple hospitals are involved, they act as one organization using the same roles. Lovis advocates distributed RBAC management, close to the location of data access, because the users involved frequently change roles. Profiles grant access to *coarse-grained* objects such as system components and applications. They are further constrained by medical service and location, i.e. a nurse can only access patient information from a computer in her own ward. Lovis et al. have implemented an emergency override procedure that allows annotated access to otherwise restricted information. This procedure is also used to fine-tune the implementation, because audit logs have shown that the emergency procedure is used very often.
- Users often claim to have very specific tasks, which cannot be modeled by generic roles or policies [32]. Adherence to the RBAC model will result in a myriad of roles unless an organizational policy to conform to generic roles is defined and rigorously implemented. However, care should be taken that the result does not interfere with good clinical care [33].
- Variations in access policies are quite common. Bakker [32] has described variations between departments, where physicians with similar roles have different authorizations. This can easily be extrapolated to organizations.

Inter-organizational connections add additional issues:

- Inter-organizational connections also require definitions of RBAC roles and profiles that do not exist in the internal organization, e.g. a hospital granting GPs access to the hospital's EHR should also implement a GP role and a profile for GPs, although there are no GPs working inside the hospital. The situation can be even more complex in practice, as individuals may have roles that combine policies from multiple issuing authorities, for example by an employing healthcare organization and a national health system and a professional body. Ensuring that none of these policies contain conflicting rules/constraints may at times be challenging.
- The RBAC model also implies that communicating parties agree on the definition and meaning of the role and profiles, i.e. a healthcare worker with a specific role has access to similar kinds of information in various systems. Extending this to all systems requires the existence of a universally applicable model of role definitions that is adopted by all health organizations.

The American standard ASTM E1986–98 [34] has defined such a list, based on the American definition of roles. ISO DTS 21298 defines a similar set of structural and functional roles and refers to the International Labour Organisation [35]. Since both do not give definitions or descriptions of each role and its policies, the only comparison can be done on the name of the role, yet from the standards it is not clear that, e.g. the term "nurse" and the policies for the role of nurse are identical or at least comparable between ASTM and ISO.

Allowing patients to restrict access to their information reveals several discrepancies:

- If resources in RBAC models are defined as tasks that align to business processes, the patient cannot provide proper access restrictions as he is concerned with the information that might be accessed by several tasks.
- If resources in RBAC models are defined as information objects, i.e. 'medication', access rules are defined for these objects, where patients would like to define (grant or deny) access to specific instances of these objects, e.g. allow access to 'medication' while restricting access to 'medication' instances from the PSYCH-EMR.
- If a patient can grant or deny access to a single person, rather than a role, a universally unique user ID, rather than an ID issued by the organization, is necessary to implement this restriction, since this person might not be a member of the organization at the time the access restriction is issued. It would also imply that to consistently deny a person access to the entire virtual EHR of the patient, this restriction has to be communicated to all other systems, preferably by a distributed component. In other words: role-based access needs to be extended with access rules for single individuals, thus increasing complexity.

In summary, ISO TS 22600 defines a structural and architectural approach to representing all of the information that may be required in policies in a standardized form. However, a complementary challenge is how to enumerate the classes of user, classes of EHR and kinds of constraint rule *interoperability* such that specific instances of policies can be recognized and applied consistently in diverse settings. There must be semantic interoperability regarding authorization roles and their respective profiles to allow consistent authorized access of information to external parties. The CEN 13606-4 standard recognizes this. A true global consensus on RBAC roles

and policies seems very unrealistic, since it would affect the local organization of roles and might even be incompatible with national definitions. However, since the broad role concepts of "doctor" and "physiotherapist" are virtually identical throughout the world, we think that this consensus can reach a workable level. Also, all roles and policies should be implemented in all systems, even if they are not applied to internal users, to allow properly defined authorized access to external users.

As we have stated before, RBAC management is already very complex on a small scale and will quickly increase in complexity on a regional, national and global scale.

Therefore, we believe that more research needs to be done to properly define an authorization service that has the flexibility of RBAC but can overcome its flaws. This research should take the proposed environment as a starting point to consider the described issues. ISO TS 22600 has made considerable steps in that direction.

To allow patients full control over the access restrictions to their virtual EHR, the discrepancies between tasks versus roles and objects versus instances need to be addressed. A universally unique ID for healthcare users would also be necessary.

4.2.2. *Confidentiality*

Related questions:

10. If PSYCH-EMR information is stored in the GP-EMR system, how can the PSYCH-EMR system be informed of possible confidentiality breaches?

As stated in the requirements confidentiality requires that proof can be given that the information has not been made available or disclosed to unauthorized entities, whether persons or systems. This can be implemented in two ways: either information is tagged with metadata about the confidentiality status or confidentiality is enforced by access rules.

Either way, if information is shared with authorized persons (e.g. the discharge information in STEP 2), there is no possibility for the sending system to verify whether this confidentiality has been breached (e.g. when the information is passed on to third parties). Since it can be argued that it is now the responsibility of the receiving party, there should be a legal and/or technical framework that regulates this issue.

This implies that access profiles regarding the information should be added to the confidentiality tags and functionality in the receiving system should be present to execute these profiles. This requires that both systems use compatible profiles. When RBAC primarily focuses on permissions on tasks, this should be solved in a different way.

Using access rules to enforce confidentiality relies on the audit logs to verify that confidentiality has not been breached. This option also cannot verify if confidentiality breaches occur at the receiving end.

Related questions:

3. Should all information (always) be available unless restricted by the patient or legislation or should information only be available if legislation and patient consent permit? At what level of granularity should the patient give his consent? Are there distinctions in types of situations? Should delegation be allowed, e.g. only if the GP thinks it is relevant?
4. Following from the previous question: Is it legally acceptable to ask for patient consent for access by unknown health-related external parties? I.e. *any* doctor or *any* nurse versus a specific, named person.
12. What happens with this type of restriction if the patient moves from one GP to another?
13. When an emergency override is necessary, what access restrictions have still to be obeyed?
14. Can the right of the patient to delete information from his EHR be sufficiently covered by a total access restriction?

4.2.3. *Patient consent*

Patients can allow or deny sharing their information with other healthcare workers. Consent must be either implied or explicitly given before the act of sharing.

4.2.3.1. Implicit or explicit. Patients either implicitly or explicitly consent to sharing information. This covers not only the exchange of information, but also the access to information by others than the original creators. Implicit consent assumes patients to have consented unless they specifically state otherwise. This is also referred to as opt-out. Explicit consent, or opt-in, is the reverse, where access to the information is prohibited unless the patient has given consent [36].

There is much debate about consent management. Studies have proven [37,38] that the opt-out model results in simplified management and a higher number of included patients. This is also the proposal of the English NHS [37–43].

Advocates of the opt-in model state that it is the only guarantee to preserve patient privacy, which is considered to be more important than simplified management. The Royal College of General Practitioners [44] advises an opt-in approach.

Dutch legislation is in line with the EU directives, that the patient's agreement with the transfer of his health information should be laid down in an explicit consent [45]. However, in an attempt to keep the work process manageable but still compliant with legislation, the Dutch WGBO[4] law [17] defines a list of standard situations that assume implicit patient consent as well as a list of situations in which patient consent should be actively requested. Some examples: implicit consent is assumed for granting access to other healthcare providers participating in the care process and for access in

[4] This law defines the rights of patients and healthcare professionals regarding health care.

emergency situations; explicit consent is necessary for granting access for a new episode of care or for research purposes.

Either way, managing patient consent within an organization is already complex. It adds an extra complexity level to the already complex RBAC model.

Implementation of patient consent is necessary for the purpose of automatic determination whether sharing of information is legally allowed. In Coiera and Clarke this is described as the gatekeeper system [36]. The system is able to remind the user that consent is denied or should be actively requested.

The focus of patient consent, as it is presently being tackled in e-Health programs, is on access to information of a particular kind and for a particular purpose, while the RBAC model can focus on access to information or to operations on information. A discrepancy is possible between the set of information the patient consents to share and the set of information the user has access to. An example to illustrate this: a patient is treated in a clinic for sexually transmitted diseases (STD) and has given consent to access his medication records created in the STD-clinic, but only to the physicians in the STD-clinic. His treating physician in a general hospital, who is authorized according to his RBAC profile, to review the medication records, should be denied access to the medication records in the STD-clinic.

4.2.3.2. *Internal versus external*. In Step 2 of the scenario, the patient restricts (partial) sharing of his psychiatric information to his GP only. This not only supports the argument that all available roles should be defined in a system, even if they are not fulfilled by actual organization members, but also raises the issue of the extent of patient consent. If the consent is fully specified for the internal situation of an organization, there might still be open issues when external (health-related) parties access the information, as shown by the STD-example above.

On occasions, the patient's consent to information disclosure needs to be captured explicitly, with a signature, for example to disclose health data to an employer or insurer. In such situations, paper systems are still usually used, but in future, a form of digitally signed policy is more likely.

A more complex challenge for both the STD and the psychiatric scenarios above is the implicit assumption that the whole record created within one care domain can be protected from access by all other clinical domains. Whilst this could probably be done in theory, the challenges for patient safety have not yet been adequately understood: will medical errors arise if vital information from those ring-fenced care settings is not provided to others treating the patient? A common proposed approach to this is to include within the overall policy framework some kind of override mechanism.

Due to the fact that an EMR system is currently designed for usage within an organization, no definitions are available about emergency overrides by external parties, e.g. an ER-physician from another hospital.

A very probable situation where the patient consents to the requesting party during an encounter, rather than beforehand, has also not been covered. This could result in a delay in therapy when the patient has to contact the originating party, unless the current treating physician is allowed to use an emergency override procedure to access the information.

The fourth part of the EN 13606 standard defines emergency procedures to override access restrictions [9]. The standard does not specify in detail which access restrictions should be overridden.

The EU directive Articles 7 and 8 (EU 95/46/EC) [45] demand that the patient's agreement with the transfer of his health information should be laid down in an explicit consent [46], unless he is incapable of doing so in situations where it is of vital interest to share the data with a health professional. EU directives also imply that the organization sharing the information is responsible and should verify that the requesting party is indeed a qualified health professional. This implies that information is shared between natural persons. It is unclear what the legal implications are when patients consent to generic roles, which can be fulfilled by hitherto unknown persons.

In the context of the Dutch new ICT infrastructure all health care providers (both persons and organizations) will be uniquely identified through a so-called UZI-number [20]. Organizations can use this number to access external information of a patient. However, when the UZI-number of the organization is used the external system cannot properly identify the actual person that requests the information and is therefore unable to respect the patient's wishes. Even when transferring along with the information the responsibility to honor the security profiles and patient consent to the requesting organization, full identification of the requesting person and his role is necessary.

4.2.4. *Relevancy*

Related questions:

3. Should all information (always) be available unless restricted by the patient or legislation or should information only be available if legislation and patient consent permit? At what level of granularity should the patient give his consent? Are there distinctions in types of situations? Should delegation be allowed, e.g. only if the GP thinks it is relevant?
9. Dutch legislation [WGBO] only allows access to "need to know" information. Should systems be capable of deducing what information a GP needs to know, and if so, how can this be achieved?
16. Should the dermatologist be able to request all available information or only *what* is relevant?
17. Should the dermatologist be able to access all available information of the patient (present in various systems) at any point in time or only *when* relevant?
24. Should the dermatologist be allowed to access the patient information after the discharge letter is created?

Dutch legislation [17] states that a healthcare worker who uses an EHR should not access information unless it is relevant for the management of the patient. Similar legislation can be found in other countries. The General Medical Council in

the UK states in its regulations for Good Medical Practice [47] that physicians should not disclose information about their patients except in specific cases.

In the past, the owner of the data could raise questions about the relevancy of a request, but in an automated exchange, this is not possible. New Dutch legislation is likely to enforce GPs to make their data available through the Dutch infrastructure defined by NICTIZ (Dutch Institute for ICT in Health care) [48]. So the burden of determining relevancy is now completely with the requestor of the information.

This requires a clear definition of both the information that is allowed to be shared and the situations in which sharing is allowed, i.e. an unambiguous definition of 'relevancy' that pertains to both content (what is relevant) and time (when is it relevant). Yet, relevancy is a very ambiguous concept that is also highly dependent on the context and therefore very hard to define in such a way that implementation in a software system is feasible. This is recognized by Lovis et al. [31]. They have solved the issue by requiring a well-defined care relationship between user and patient and by relaxing the time constraints. This approach is less suitable in a cross-organizational situation since the requested system cannot reliably verify a suitable relationship between the intended external recipient and the patient.

Conversely, relevancy prevents data overload: allowing a user to view all available data can result in losing the overview and the capability to make an informed decision. Striving for the display of only relevant information becomes a pragmatic goal.

The ideals of relevancy include not only the right to know (if a suitable care or consented relationship exists with the patient) but also the need to know (granting access only to the health record information that is needed by the requestor to perform relevant tasks). However, determining the latter is almost impossible in advance for any individual patient and any individual access scenario (i.e. when data are committed to an EHR or when policies are being defined). In our opinion relevancy of information can rarely be defined a priori and can only be implemented in a few well-defined, unambiguous situations. It is more pragmatic to allow access when in doubt and verify possible abuse afterwards (e.g. through the audit log).

Prevention of data overload could be implemented by a layered structure that shows an increasingly detailed view of the available information. On the top level, only the different types of information are available. A second level shows summaries, while the lowest level shows detailed data. This allows users to drill down to the most interesting, i.e. most relevant, data while noticing the availability of other information types.

4.2.5. *Ownership of information*

The current practice of storing both locally and externally created information in the same system requires the identification of the owner and the origin of the information. The origin of the information is often defined as the location where the information is created, i.e. entered into the system.

In general, the owner is defined as the creator of the information. Establishing the owner of the information is necessary for several purposes: the owner is responsible for the availability of the information [46], for the accuracy of the information[5] [17] and for protection against unauthorized access [49]. The owner is also held responsible in legal disputes. Note that the 'owner' can refer to the person responsible for the information or to the organization storing the information.

Related questions:

6a. How can the GP trust the information?
19. If the information is assumed to remain with the owner, who is the owner? The person ordering the information or the person generating the information? Or both?
23. Should the GP be able to differentiate between the various owners of the information included in the message?

Ownership is closely linked to origin. Information that originates from an external system has a different owner and should be processed differently. Blobel has argued that data should be traceable to its origin and that data should be kept only at the origin [50] to avoid redundancy and to maintain integrity.

Clear definition of the ownership of data is necessary to resolve the question whether the lab is the owner of the lab results or the physician ordering them, because it is directly relevant to the issues mentioned above. Only the lab can send corrected test results (maintain accuracy), but the dermatologist is responsible for acting on them, and for including the correct lab results in his discharge letter.

Definition of ownership is also important to resolve the situation where a patient chooses a different GP and all of his information has to be transferred. The new GP has not created the information, but on using it (and thereby trusting it), he is also responsible for assessing the accuracy of that information.

There is some ambiguity in the above definitions. The Dutch WGBO law states that the creator of the information is the owner of the information [17]. NICTIZ however, differentiates between the author and the "manager" of the data [51]. The manager has not created the data but incorporates the information in his/her own system and assumes responsibility. An example to clarify this is the Dutch situation where many regions have a central General Practice locum service (HAP) where GPs are on duty in the evening or weekend. When a patient is treated at this HAP, a summary message is sent to the regular GP who incorporates the information in his local system. From that moment on, the source of the data is the regular GP's system. The GP is also the manager, while the locum remains the creator. This distinction solves the problem of transferring ownership. However, it also implies that data is not strictly stored at the place of creation (i.e. the system of the locum).

[5] Initial accuracy is the creator's responsibility, who is also the owner at the time. After transfer of the ownership, the new owner might be held responsible for the accuracy.

Table 1 – Definitions of ownership

	Definition
Source	Location where the data are stored
Origin	Location where the data are created
Manager	Person/entity responsible for the data (provide, protect)
Author	Person/entity responsible for the content of the information
Creator	Person generating the data and entering it in the system

This led us to define the concepts of owner and origin more precisely and make a distinction between the creation and the management of the information.

The discussion above shows that the definition of 'owner' is very ambiguous. Legislation in different European countries also differs on defining ownership of the data in an EHR, which further increases the ambiguity. We therefore propose to avoid using the term "owner" and want to introduce the following terms:

- "creator" for the person generating the data and/or entering the data into the system;
- "author" for the person or entity responsible for the content of the information;
- "manager" for the person or entity responsible for the management, provision and protection of the information.

Similarly, the source is the location of the information that is managed by the manager, while the origin is the location where the information is created. We assume that in most cases, both persons and locations, respectively, are the same (Table 1).

Correct definition of manager, author and source are directly related to data integrity and non-repudiation (both receiving and sending) requirements of the EHR. It is possible to define business and legal policies for situations where the manager and/or author of information changes (e.g. when a patient changes to a different GP) or where information from several authors is exchanged with third parties (e.g. the discharge letter from the dermatologist to the GP that includes test results). In the latter case, there should be clear definitions on how the external information should be incorporated to avoid conflicting authorization policies and outdated information.

4.2.6. *Infrastructure*

Information in a system is subject to changes caused by corrections such as the lab test rerun described in the scenario (STEP 11). Audit logs capture the correction of data, while version management ensures that each value (i.e. version) of the data is clearly identified.

Exchanging information with external systems requires a notification mechanism to inform (external) recipients of the correction.

4.2.6.1. *Version management*. Version management is necessary to uniquely identify a specific value of a specific concept at a specific point in time, for example in recreating a specific situation during an audit.

This implies that the version management of the system is capable of retrieving not only the most recent version, but also the version that was previously seen by the user. This in turn implies that queries to external systems should include parameters that can be used by the system where the data resides to retrieve the correct version.

Version management in distributed systems requires the use of unique IDs for each version of all relevant objects, whether persons, devices or data structures and for a common identifier for all versions of the same object. This ID is separate of, but connected to the object's ID, which should also be available in each version.

Related questions:

15. How should the query be propagated to all available systems in the area?
20. How should the dermatologist be informed of the corrected information?
22. Should the dermatologist be able to view both the current (i.e. new) and the old (i.e. those sent before the recalibration) values?
25. What action should be taken if the test results were updated after the discharge letter was sent?

4.2.6.2. *Notifications*. When information is updated, all relevant parties should be notified, including external parties, which implies that the system providing the information registers the requesting parties in such a way that this information can be accessed and used.

In general, notifications are used at a lower OSI level to handle asynchronous processes such as transactions. In this section, we discuss notifications at an application level where they are used to inform interested parties of updates. Notification is not only used to secure data integrity, but also to inform health care professionals that they have (possibly) based their decision on possibly invalid data.

If we look at the scenario where the dermatologist writes a discharge letter based on the test results available and the lab sends corrected test results, two situations can be distinguished:

- Only decisions based on external information are described in a document that is sent to a third party. In this case the corrected information is not included, but might result in different decisions and therefore an updated document.
- External information is incorporated in a document that is sent to a third party (e.g. test results are included in the discharge letter sent by the dermatologist to the GP). In this case not only the dermatologist should be notified, but also the GP, regardless of a possible update of the discharge letter with a different decision.

Both cases imply that the sending system should have a registration of previously requested information and the identification of the requesters.

This also requires the registration of all external information that is incorporated in information that is sent to third parties. If the corrected information matches the external information, the user should be informed and an updated version of the including information should be sent to the third party. A question arises if the user should be informed of the corrected information if he has not seen the previous version, since he has not acted upon the invalidated information. Whatever the outcome of this question is, it always requires a registration of which user has retrieved what information. Another approach is to only send notifications of corrected information when the information has been queried before. This implies that an EHR system not only registers the senders of the queries, but also tags the information with properties to determine if previous versions of the updated information have been queried.

In all cases, an interoperable notification framework is necessary.

It is not clear where to locate this notification registry. Although this information is also logged in the audit log (see Section 4.2.7), we believe it is against the purpose of the audit log to be actively used by the EMR system. This would result in a very convoluted link between the EMR system and the audit log. Setting up such a registry in the EMR results in duplication of information, which is also not a desirable situation.

There is also concern that the automatic notification of every EMR data update to every previous potential recipient is a logistic nightmare, and may result in unwanted disclosures (if, for example, previous recipients are no longer caring for the patient).

4.2.7. *Audit*

Related question:

21. Should the system be able to log the fact that the dermatologist has actually seen the (updated) information?

The audit log should document all actions performed on the information and the users performing those actions, to enable the recreation of the state of the past.[6] There are various reasons for this functionality:

- It can be used for security purposes, i.e. to monitor the access and possible misuse of the system, preferably in real-time.
- It can be used for review purposes; EN 13606-4 [9] specifies how patients might be provided access to audit log information to review access to their EHR. This implies that the log information should be converted to something that is understandable by the patients. Security policies should be applied for the auditors to avoid security breaches. The Dutch infrastructure of NICTIZ will have this functionality implemented.
- It can be used in legal disputes to verify claims about what information was available and whether it was accessed. This can also be used for non-repudiation issues.
- It can be used to complement access control to update and fine-tune RBAC policies as well as to verify the relevancy of the accessed information.

The logging should not be restricted to the information handled, but should include all events and state changes. This results in a huge volume of data that should be stored for (possible) future reference, while it should also be possible to process the log in real-time to detect intrusion or unusual behavior. The use of an emergency override mechanism to access otherwise prohibited health record information is an example of this.

Due to the distributed nature of information retrieval in the context of the scenario of this article, information and audit logs are also distributed. In order to track the full details of an action (e.g. a query and its results) many systems and subsequently many audit logs are involved. This trace across systems requires interoperable audit logs to allow unambiguous processing. To faithfully trace a specific event across systems all events need to be uniquely identified through a generic identifier, either by the same identifier throughout all systems or by mapping the external identifier to an internal identifier. Furthermore, a synchronized timestamp is necessary to order the events and functionality is required to retrieve a specific version to allow faithful recreation of a past situation.

Although the interoperability is recognized by the EN 13606-4 standard [9], no interoperable audit logs currently exist. The IETF RFC 3881 [52] provides a draft specification. ISO and CEN are currently working on a standard for a common framework for interoperable audit logs based on this RFC.

The requirement to maintain an audit log can be implemented in an audit service.

Audit logs should be separated from the EMR system, with a clear definition of the connection between the two to maintain data integrity on both components. This allows the use of the audit log as described above, without placing a load on the EMR system itself. In addition, different authorization profiles apply to the audit log, those of audit managers, not of healthcare professionals. Logging actual data in the audit log should be avoided, because of the security issue. Otherwise restricted information would then become available to non-healthcare professionals (e.g. audit managers). By logging a reference to the information, rather than the actual information, in a separate system both requirements can be met. The audit logs relating to emergency access to the EHR, which often overrides usual access policies, should be reviewed extra carefully to ensure that such privileges are not abused. A separate log event with details of the emergency action could be created in a separate, highly restricted service, with a minimal reference in the regular log. The latter would be sufficient for routine

[6] Note that clinical audit, which we define as the auditing of the medical information for the purpose of improving health care, is beyond the scope of this article.

audit procedures, while the former can only be accessed after proper authorization. Emergency overrides should be investigated to determine the necessity. This requires recording of the situation that triggered the override, information that is often not recorded.

In some countries, patients have the right to have data removed from their health record. This would imply that the audit log contains references to non-existing data. Provisions are needed to cope with this situation. Data removal by patients also poses problems when trying to recreate past situations, since the data is not available any more and should be marked as such.

Accessing the audit log itself should also be logged for the same security reasons, although with the theoretical risk of creating audit logs ad infinitum.

Storage requirements of an audit log can be considerable which might lead to policies to keep logs only for a defined time span. Since the cost of storage devices is continually decreasing, the only reason for a policy of purging audit logs should be legal requirements.

Use of the information in the audit log might have medico-legal implications. Based on the logs it might be deduced that certain information was accessed without imminent need. It might even be deduced that information was *not* accessed when access could be expected. Although it is obvious that unauthorized access should be investigated, what are the implications of the *absence* of access?

This discussion shows the importance of a clear definition of the purpose and content of an audit log as well as a clear model of the connection between EMR system and audit log.

Such a definition would dictate the kinds of information that an audit log should contain, and the ways it should be monitored.

4.2.8. *Archiving*

Related question:

26. What should be done with the data after legal data retention time has passed?

Archiving is the mechanism of moving data out of the active system into locations, which are less immediately available, usually for storage logistics and performance reasons. Wherever possible, archived data should be technology-independent so that future users do not have to depend on obsolete technology from the past.

Dutch legislation [17] defines that medical data should be kept for at least 15 years. Longer periods are allowed if it is essential to the health of the patient or even to the health of his relatives (e.g. hereditary diseases). Many countries require longer data retention periods for liability issues.

This implies that systems should be able to retain this information for at least the period of the legal data retention. For such a long time, the issue of the ownership is even more important, but it also complicates queries, because older information is usually less relevant in everyday treatment.

Not only the data, but also the authorization profiles pertaining to that data should be preserved. Roles and profiles will evolve over time and reliably answering auditing questions can only be done when the exact context of the question is recreated, including the authorization profiles at that time. This in turn implies that version management on roles and profiles is necessary.

Records spanning a lifetime with multiple care providers both in parallel and in succession and restricted by both organizational and patient restrictions might lead to data that should be retained due to its age and/or importance but is not available due to its authorization profiles.

In paper-based records older information is often manually summarized, e.g. after each episode of care, reducing the size and superfluous details in the data to be interpreted in future episodes of care. This does not imply that the detailed data is removed. The electronic record could also provide summaries while retaining the detailed data. There is no standard yet that deals with the action of summarizing electronic data, either manually or automatically. There are also no guidelines on how to assess, which information is still relevant for the patient's situation after the retention time, who should make this decision, and on what criteria this decision should be based.

Dutch legislation allows patients to ask for removal of data from their record. This not only overrules the data retention policy, but also contradicts the ISO 18308 requirement that data should not be deleted. Data retention policies and lifelong health records are also contradicting since the latter requires all information available during the lifetime of the patient, while the former requires destruction of information older than a certain age.

Finally, documentation obligations of organizations (e.g. for health statistics or liability requirements) conflict with the patient's right to remove data. One approach could be to allow the patient to hide direct access to information using a mechanism outside the regular access control mechanism, e.g. by using encryption with the patient's key (thus requiring the patient's cooperation for decryption), whilst permitting it to be included in population queries that are needed for statutory reporting.

Data retention policies and lifelong health records require long-term storage. This means that data should still be accessible long after the creation date. This implies storage in future proof formats and/or keeping old software and hardware components around to be able to access it. It puts considerable strain on the various implementations.

Shabo [12,13] recognizes these problems and describes a solution where health records are maintained at health record banks equivalent to financial banks. He suggested that the health record banks should be funded indirectly through the health insurance plans of the patients opening an account. More recently, a number of large commercial players have launched personal health record systems. Although this solution moves the burden of long time storage to a central health bank or central commercial server, the problem of future accessibility will remain.

The question arises to which extend should past situations be recreatable. This period could be much shorter than

the retention policies. In special cases, manual procedures are probably the most cost-effective.

4.3. *Functionality and implementation of the central service*

The central service as depicted in Fig. 1 plays a vital role in the described environment. Literature shows two approaches for implementation of such service, which we will describe here:

- The medical information is stored in a central system: a repository.
- Pointers to the medical information are stored in a central system: an index service.

The repository can vary from a set of (pointers to) attested documents about specific aspects of the medical history of patients [53] to a full record of all available information [12,13] and anything in between. A centralized index service or registry is currently implemented by the Dutch NICTIZ [54]. They are compared and contrasted here in relation to the issues described before.

4.3.1. *An lvEHR repository*

If the central service were implemented as a repository, it could contain documents covering only information that is considered of relevance for properly dealing with future health issues, like discharge letters and medication lists. Its advantage is the single point of retrieval of attested documents. No discussion on the quality and no extended searches of connected systems for information are necessary. The authorization issues still apply since only authorized persons should have access to the documents. The confidentiality issue can be avoided by adding to the repository only documents that have the same confidentiality tags. Patient consent can be handled easily because it is known beforehand what will be added to the repository and what user roles are allowed to review the information. Relevancy is handled by a strict definition of the content of the documents. No corrections are necessary (i.e. documents are aggregated/summarized from underlying information).

The documents are summaries or aggregations of information on different levels of granularity, with the actual information still residing in the various connected systems. The documents contain patient identifiable information and therefore become a privacy risk. The documents should be tagged with IDs from the authors and source, for proper definition of the ownership and responsibility of the data. Since the repository does not necessarily hold all available information, additional information either cannot be retrieved or should be retrieved through an indexing service. Connected systems are still responsible for the information and contain a duplicate of the information that is contributed to the repository.

An alternative is the PING architecture described by Simons et al. [55]. This is a centralized, cross-organizational repository of documents where the patient can define fine-grained control in delegating access to portions of his EHR.

Shabo goes one step further and claims that lifelong patient-centered EHRs can only exist when storage and maintenance is transferred to a health record bank (similar to a financial bank). The health record bank can transfer a copy of the EHR (all or in part) to a specified health care provider after patient consent. During treatment, the health care provider can add information to the EHR, which is transferred back to the health record bank after discharge. Intermittent updates of the original EHR in the health record bank are necessary when treatment takes a considerable amount of time, e.g. in a chronic condition.

The health record bank takes care of the archiving problems and the authorization, although the authorization issues remain. All available information on the patient is in his EHR account, so an indexing service for additional information is not necessary. It cannot be avoided that patients open several accounts to hide sensitive information. There is no possibility for the health record bank to verify confidentiality breaches when the EHR is transferred to a health care provider, unless all audit information pertaining to the specific record is also added to the account. It is unclear if patients can consent to sharing only part of the information or can restrict access to specific persons or roles. RBAC models are likely to differ between care providers so it is not clear which models (care provider or patient) will take precedence. It is also not clear how emergency overrides should happen when a patient has not consented to transfer information to a specific care provider and/or the EHR is not (yet) transferred to the care provider. When a patient receives treatment at different locations, careful synchronization of information is necessary to avoid conflicting treatments due to lack of information.

In both lvEHR and Health Bank repository-types, ownership is a problem: as stated before according to Dutch law, the person entering the information in the health record is the owner of the information and is responsible for its accuracy and completeness. Other countries have different legislation, making a single approach across countries very difficult.

The repository of attested documents relies on the connected systems to keep the respective documents up-to-date and complete. The same goes for the health record bank.

Authorized access in hospitals is often linked to a specific relationship with the patient. It will be difficult for the health record bank to verify such a relationship.

EHR systems storing all incoming information (as most current systems do) will be superfluous and may have no legal status in the environment of health record banks.

4.3.2. *Index service*

An alternative implementation of the central service would be as an index service. An index service does not store the information on a patient, but merely index pointers to the actual source of the information.

Various indexes are possible:

- A *passive index*, which acts as a telephone directory: the index returns locations of the information which are in turn queried for the actual information;
- An *active index*, which acts like a broker: the query is addressed to the index service. The index service in turn propagates the query only to systems that have registered the information in the index that matches the query. Finally, the combined result set is returned to the query originator.

The Act Reference Registry (ARR) of HL7 as implemented by NICTIZ is an example of an active index service.

To avoid returning unallowed information a broker system should be able to either pass on the user ID and RBAC profiles that are attached to the queries, or be able to match the profiles to the RBAC information that is added to the registered information.

Registered index pointers should only be visible to users with appropriate authorization. Querying would be most efficient when only the systems are addressed that will return information for that user. This implies that a broker system is capable of executing RBAC policies. However, only the originating system is responsible for filtering unauthorized access and updating the local RBAC profiles. To execute the RBAC policy by the broker would imply that the entire set of policies of all registered systems should be duplicated in the broker system. This seems to be an unacceptable situation.

A separate complicating factor is the granularity of the information index. Coarse-grained indices can be fewer and so more manageable, but may lack sufficient details for RBAC policies and/or patient restrictions.

Legislation might deny storing privileged information in a central location. This effectively reduces the options to a passive or active index service.

5. Discussion

5.1. *Discussion of the approach*

Our method for identifying the relevant issues and their elaboration is similar to the first two steps of the methodology of the HL7 Message Development Framework (MDF) [56] for defining messages.

In very global steps, the HL7 approach starts from a realistic scenario or storyboard through the definition of use case models and interaction models, using UML modeling tools, to specific messages. A scenario is a more informal description of the events, which is better suited for capturing the time-sequence of events than the more formal use case models. It also hides details of the Actors involved in the interaction.

Following the HL7 MDF approach allowed us to focus on a specific scenario, which simplified the definition of the questions. Only one scenario was defined, but it represents a generic, privacy sensitive situation. Each institution or department can be replaced by a different equivalent where sensitive or general information is generated, without changing the overall issues.

By focusing on real life processes and asking simple questions, we believe the essence of an issue is much better clarified than defining an abstract goal that should be met. An example: RBAC is considered to be a good solution to define access rules for users and its implementation would meet the goal of secure systems. However, a question like “how could a dermatologist request information from a GP system” shows that implementing RBAC in both systems is not sufficient to solve the problem raised by this question.

Defining questions is an endless process; there will always be more questions to ask. By grouping the questions in issues, which are then mapped onto the topics of various standards, it can be argued that the question set is sufficiently complete when all issues cover all topics.

Table 2 – Mapping of issues to EHR requirements

Issue	EHR requirement
Authorized access	Authentication/authorization/patient access
Ownership	Integrity (of data)/non-repudiation
Ownership	Author responsibility
Confidentiality	Confidentiality
Patient consent	Consent
Relevancy	
Audit	Audit/non-repudiation
Version management	Audit
Notifications	
Archiving	Archiving

Table 2 shows that the topics of data integrity and non-repudiation have not been mapped to a separate issue but depend on proper identification of the author/manager of a record contribution and appropriate audit logs. Patient access is also not covered separately, because it would make the scenario even longer and more complex, and because a patient can be seen as a user with a special role and profile and therefore not different from other users such as health care providers.

The only mention of relevancy in the ISO definition is the requirement that only data relevant to the care process should be recorded. It is not mentioned as a separate requirement. Version management also has no corresponding requirement but the discussion shows that it is a vital pre-requisite for auditing. The issue raised by notifications has not been covered by any ISO requirement although it touches on data integrity.

Unraveling the issues gives insight into the problems that can arise. Literature review has shown various studies into these problems and possible solutions. However, most reviews focus either on a single issue or on a limited environment. We have found that the issues influence each other and a solution for one issue can cause additional problems for another issue.

Fig. 3 gives an overview of the relations between the issues. We see that all issues are interrelated, but that authorization, patient consent and audit play a pivotal role.

We used an approach where we discussed and studied the issues in the order presented in this paper. Information and possible solutions from an earlier issue were taken into account into later issues. This sequence was repeated several times, while taking into account all the information that was revealed in earlier rounds. When no significant changes in the information occurred, the issue was regarded as “closed”.

5.2. *Discussion of the results*

The lvEHR can only become reality when there is a truly secured network of EMR systems that contain information about the patient. Thus, in theory each patient has a different network, i.e. comprised of different EMR systems, which is not fixed over time since it can grow when the patient visits other healthcare providers/organizations. As described before,

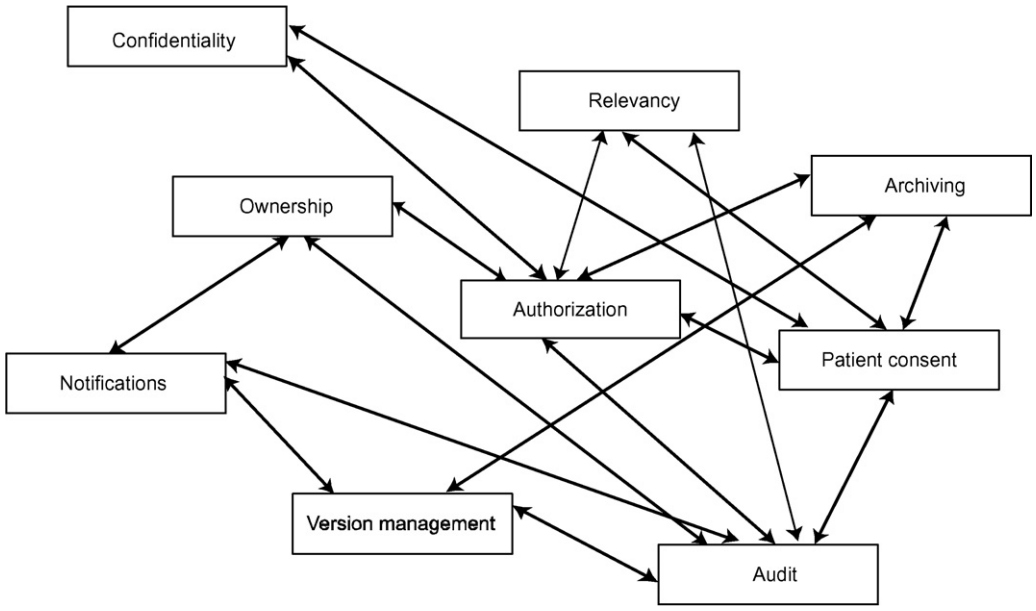

Fig. 3 – Relations between issues.

there are some issues to be solved before these networks will be in operation.

RBAC as described in this article is considered more suitable for simpler environments. More complex environments such as health would benefit from policy-driven RBAC, as described in the Privilege Management and Access Control (PMAC) framework. This ISO specification [27,28] extends RBAC by defining a framework to represent and manage computable policy agreements between parties to exchange and use information. The policy agreements specify which information to exchange and under which security-related circumstances. Although these policy agreements simplify the problem by focusing each policy on one type of exchange under a specific set of well-defined circumstances between parties, the level of overall complexity rises with the number of agreements since these each need to be maintained and updated to reflect current changes and also need to be in accordance with each other's other policies to avoid conflicting policy agreements.

We believe that PMAC will certainly solve most problems mentioned in this article at a very restricted level of cooperation between a few parties. Expanding the cooperation to include more information and more parties at a greater regional and national scale will require a lot of resources to keep all policies in line.

Since policies agreements are defined up front and go through a standardized definition process, there is no possibility of securely exchanging new types of information that have not yet been agreed on.

We think the current paradigm shift from paper-based to the current EMR systems with their local storage of all incoming information is not sufficient to transparently incorporate a virtual lifelong patient record that spans multiple independent healthcare organizations. The shift should be extended to move towards an infrastructure that resembles the environment described in this article with a focus on interoperability and sharing of information in privacy enhanced way.

True interoperability of systems pertaining to an lvEHR can only be guaranteed if the privacy issues mentioned before are solved. It also entails not only semantic interoperability of the information, but also of the authorization roles and profiles as well as all other aspects pertaining to networked EMR systems.

We have focused on the security issues in an environment of interconnected EMRs/EHRs. Currently the notion of a Personal Health Record (PHR) is elaborated. The PHR is seen as a solution to guard against identity theft and unauthorized access since the only way that data can travel is from an EMR into the PHR. With the patient in control of the PHR, data filtering is performed by the patient (some data will be included, other data not). Access control is also executed by the patient, who can decide during a consultation whether to share existing data or not with the consulted physician. We have not analyzed the security issues related to the PHR in detail. We see that besides the security issues, other aspects like the impact of unshared data on patient care, etc. play an important role in the further development of the PHR. Such issues are outside the scope of this manuscript.

A promising technology for implementation of the described environment would be peer-to-peer (P2P) networking [57], where all the participating systems share the computing and storage resources to fulfill the demand. This would result in a truly transparent EHR system. However, since P2P also entails replication of (parts) of data on several nodes, it raises security issues along the lines, described in this article. More research into that area is necessary before it can be useful.

6. Conclusion

In the age-old tradition of medicine, healthcare workers have developed a way of communication that follows the lines

of a conversation. Information is specifically requested from defined other persons. Typically, even today the preferred means of communication are still telephone, fax and email. Current information retrieval as defined by HL7 and 13606 is still very much based on this practice.

The currently omnipresent search engines on the Internet that provide access to vast amounts of information have shifted the paradigm from a one-to-one exchange to a one-to-many exchange and from a tendency to store everything locally to a tendency to search and retrieve whenever the need arises. It is time to implement that paradigm in the EMR systems and the lifelong virtual patient record.

Before the described virtual lifelong patient record can become reality, more clarity has to be provided on the following:

- The implementation of the authorization model.
- The implementation of patient restrictions and patient consent in general.
- Legal and computational frameworks that protect confidentiality.
- A definition of relevancy that is legally accepted and machine interpretable.
- A future proof infrastructure that supports communication between systems in different organizations given the above restrictions.
- Archiving information for future use given the situation where people and systems are changed frequently in the course of the life of the patient.

Technical solutions need not and will not replace *all* personal communication. The EHR as a shared source of information for professional discussion can pragmatically solve some of the issues (such as the relevancy issue) discussed in this article.

Summary points

What was already known

- Electronic health record systems need to provide adequate functionality to restrict access.
- The Role-Based Access Model (RBAC) provides a sufficiently flexible method of granting access to information.
- Audit logs are necessary for various reasons such as verifying access.

What this study added

- A clear overview of what issues need to be dealt with before a lifelong virtual EHR can become reality.
- Adequate restrictions within an organization need to be extended to external parties allowed to access the patient information.
- Granting secure access to external parties requires consensus on the definition of roles and profiles across organizations.
- Audit logs need to be semantically interoperable to allow faithful reconstruction of a chain of events across systems of various organizations.

Acknowledgements

Author's contributions: H.v.d.L. is guarantor of the study. She initiated the idea for this review. H.v.d.L. and J.T. discussed the scope of the review, made an inventory of the issues and outlined the first approach. In several meetings H.v.d.L., J.T., D.K. and A.H. elaborated the issues in more detail.

H.v.d.L. wrote most of the manuscript. J.T., D.K. and A.H. critically reviewed the various versions of the manuscript. All authors approved the final version.

Appendix A

A patient, Mr. Jones, lives in a town with a large hospital, a small psychiatric institution and several general practitioners practices. These are all separate organizations with contracts to share relevant information. The patient has one specific GP whom he regularly visits and who is fully informed about the patient's medical history.

Mr. Jones has a history of depressions that once resulted in a short stay at the psychiatric institution (STEP 1). He is doing well now and his current medication prevents relapse into depression. Mr. Jones has given his permission to the psychiatric institution to respond to requests for information only from his GP (STEP 2). The psychiatric institution has sent discharge information to the GP (STEP 3). Mr. Jones has also informed his GP that he does not want his psychiatric records be disclosed to others, unless it might have serious implications on future treatment (STEP 4).

One day Mr. Jones develops a rash and consults his GP (STEP 5). The GP is unsure whether it is an allergic reaction or something else. Rather than investigating the issue himself, he decides to refer Mr. Jones to the dermatology department at the hospital (STEP 6).

The dermatologist wants to know whether there is medical information about the patient elsewhere, regarding allergies and medication. The answers of the patient are vague, so he decides to ask the GP[7,8] (STEP 7), who responds in compliance with the patients wishes (STEP 8). The dermatologist orders also a blood test and a skin allergy test to be performed at the hospital labs (STEP 9). The lab performs the tests and sends the results to the dermatologist (STEP 10). The skin test reveals a mild allergic reaction to cats and the blood test results show slightly elevated levels.

One day later the lab sends new blood test results (STEP 11). The machine turned out to be at fault and was recalibrated and the tests were redone. The new results are all within normal ranges.

[7] We assume the patient has given consent to the dermatologist to request information from the GP.

[8] It might be reasonable to assume that it is good practice to include this information in the referral letter. For this scenario we assume it has not been done.

In the next consultation, the dermatologist advises the patient to stay away from cats, and prescribes tablets to reduce the itch (STEP 12). Mr. Jones is referred back to the GP (STEP 13).

Appendix B

The entire list of questions, grouped by scenario step.

- Step 2: The patient informs the psychiatric institution to restrict sharing of PSYCH information to only the GP.
  1. How should a patient be identified reliably across organizations?
  2. How should health professionals be identified reliably across organizations? How should organizations be reliably identified?
  3. Should all information (always) be available unless restricted by the patient or legislation or should information only be available if legislation and patient consent permit? At what level of granularity should the patient give his consent? Are there distinctions in types of situations? Should delegation be allowed, e.g. only if the GP thinks it is relevant?
  4. Following from the previous question: Is it legally acceptable to ask for patient consent for access by unknown health-related external parties? I.e. *any* doctor or *any* nurse versus a specific, named person.
  5. Can the patient consent to sharing his PSYCH information while consulting his GP? How should this be implemented?
- Step 3: Information from the PSYCH-EMR is shared with the GP (in compliance with the patient's consent).
  6. How should the PSYCH-EMR system define authorization of the GP to access information in the system?
     a. How can the GP trust the information?
  7. Should all systems have authorization information for all possible users (i.e. persons requesting information)?
  8. Should all systems provide similar access for all possible users (i.e. a GP has access to the same kind of information in all systems)?
  9. Dutch legislation [WGBO] only allows access to "need to know" information. Should systems be capable of deducing what information a GP needs to know, and if so, how can this be achieved?
  10. If PSYCH-EMR information is stored in the GP-EMR system, how can the PSYCH-EMR system be informed of possible confidentiality breaches?
- Step 4/8: The patient informs the GP to not share his PSYCH information.
  11. In case the information from the PSYCH-EMR is stored in the GP-EMR system and matches a future query from an external system (e.g. the query from the DERM-EMR), should the information be passed on if patient consent permits or should external information always be excluded from a result set?
  12. What happens with this type of restriction if the patient moves from one GP to another?
  13. When an emergency override is necessary, what access restrictions have still to be obeyed?
  14. Can the right of the patient to delete information from his EHR be sufficiently covered by a total access restriction?
- Step 7: The dermatologist requests information regarding allergies and medication of the patient.
  15. How should the query be propagated to all available systems in the area?
  16. Should the dermatologist be able to request all available information or only *what* is relevant?
  17. Should the dermatologist be able to access all available information of the patient (present in various systems) at any point in time or only *when* relevant?
- Step 9: The dermatologist orders several tests.
- Step 10: The lab returns the test results.
  18. Is there a difference between the two implementations in how the dermatologist is informed?
  19. If the information is assumed to remain with the owner, who is the owner? The person ordering the information or the person generating the information? Or both?
- Step 11: The lab sends corrected test results.
  20. How should the dermatologist be informed of the corrected information?
  21. Should the system be able to log the fact that the dermatologist has actually seen the (updated) information?
- Step 12: The dermatologist sees the test results.
  22. Should the dermatologist be able to view both the current (i.e. new) and the old (i.e. those sent before the recalibration) values?
- Step 13: The dermatologist writes a discharge letter, which includes the results of the skin allergy test.
  23. Should the GP be able to differentiate between the various owners of the information included in the message?
  24. Should the dermatologist be allowed to access the patient information after the discharge letter is created?
  25. What action should be taken if the test results were updated after the discharge letter was sent?
  26. What should be done with the data after legal data retention time has passed?

REFERENCES

[1] W. Grimson, D. Berry, J. Grimson, G. Stephens, E. Felton, P. Given, R. O'Moore, Federated healthcare record server—the Synapses paradigm, Int. J. Med. Inform. 52 (1998) 3–27.
[2] J. Grimson, W. Grimson, D. Berry, G. Stephens, E. Felton, D. Kalra, P. Toussaint, O.W. Weier, A CORBA-based integration of distributed electronic healthcare records using the synapses approach, IEEE Trans. Inf. Technol. Biomed. 2 (1998) 124–138.
[3] W. Grimson, B. Jung, E.M. van Mulligen, A.M. van Ginneken, S. Pardon, P.A. Sottile, Extensions to the HISA standard—the SynEx computing environment, Methods Inf. Med. 41 (2002) 401–410.
[4] Harmonisation for the security of web technologies and applications (last accessed 2008/02/18); Available from: http://www.telecom.ntua.gr/(HARP/HARP/INSIDE/Inside.htm.
[5] HL7 (last accessed 2008/02/15); Available from: http://www.hl7.org.
[6] Health Informatics – Electronic health record communication—Part 1: Reference Model, CEN/TC251 EN13606-1, 2007.

[7] Health Informatics – Electronic health record communication—Part 2: Archetypes, CEN/TC251 EN13606-1, 2007.
[8] Health Informatics – Electronic health record communication—Part 3: Reference archetypes and term lists, CEN/TC251, EN13606-3, 2008.
[9] Health Informatics – Electronic health record communication—Part 4: Security requirements and distribution rules, CEN/TC251 EN13606-4, 2007.
[10] Health Informatics – Electronic health record communication—Part 5: Messages for exchange, CEN/TC251 prEN13606-5:2007:E, 2007.
[11] openEHR (last accessed 2008/02/15); Available from: http://www.openehr.org.
[12] A. Shabo, A global socio-economic-medico-legal model for the sustainability of longitudinal electronic health records. Part 1, Methods Inf. Med. 45 (2006) 240–245.
[13] A. Shabo, A global socio-economic-medico-legal model for the sustainability of longitudinal electronic health records. Part 2, Methods Inf. Med. 45 (2006) 498–505.
[14] M.W. Jaspers, P. Knaup, D. Schmidt, The computerized patient record: where do we stand? Methods Inf. Med. 1 (Suppl. (45)) (2006) 29–39.
[15] ANSI, ISO/TS 18308 Health Informatics—Requirements for an Electronic Health Record Architecture, ISO 2003.
[16] ANSI, ISO/TR 20514 Health informatics – Electronic health record—Definition, scope and context, ISO 2005.
[17] Wet Geneeskundige Behandel Overeenkomst (WGBO) (last accessed 2008/02/18); Available from: http://www.hulpgids.nl/wetten/wgbo.htm.
[18] R.J. Anderson, Security in Clinical Information Systems (last accessed 2008/02/18); Available from: http://www.cl.cam.ac.uk/users/rja14/policy11/policy11.html.
[19] Person Identification Service Specification (last accessed 2008/02/18); Available from: http://www.omg.org/technology/documents/formal/person_identification_service.htm.
[20] Dutch Unique Healthcare Provider Identification Register (UZI-register) (last accessed 2008/02/18); Available from: http://www.uziregister.nl/.
[21] M.J. Bittle, P. Charache, D.M. Wassilchalk, Registration-associated patient misidentification in an academic medical center: causes and corrections, J. Commun. J. Qual. Patient Saf. 33 (2007) 25–33.
[22] GS1 (last accessed 2008/02/18); Available from: http://www.gs1.org/.
[23] Registratie en informatie beroepsbeoefenaren in de zorg (last accessed 2008/02/18); Available from: http://www.ribiz.nl/.
[24] Role Based Access Control (last accessed 2008/02/18); Available from: http://www.csrc.nist.gov/rbac/.
[25] Role Based Access Control, ANSI INCITS 359-2004, February 2004.
[26] B. Blobel, R. Nordberg, J.M. Davis, P. Pharow, Modelling privilege management and access control, Int. J. Med. Inform. 75 (2006) 597–623.
[27] ISO/TS 22600-1 Health informatics—Privilege management and access control—Part 1: Overview and policy management, ISO 2006.
[28] ISO/TS 22600-2 Health informatics—Privilege management and access control—Part 2: Formal models, 2006.
[29] eXtensible Access Control Markup Language (XACML) (last accessed 2008/02/18); Available from: http://docs.oasis-open.org/xacml/cd-xacml-rbac-profile-01.pdf.
[30] M. Evered, S. Bögeholz, A case study in access control requirements for a Health Information System, in: Proceedings of the second workshop on Australasian information security, Data Mining and Web Intelligence, and Software Internationalisation—Volume 32 (2004), Australian Computer Society, Inc., Dunedin, New Zealand.
[31] C. Lovis, S. Spahni, N. Cassoni, A. Geissbuhler, Comprehensive management of the access to the electronic patient record: towards trans-institutional networks, Int. J. Med. Inform. 76 (2007) 466–470.
[32] A.R. Bakker, The evolution of Health Information Systems, security in practice and open issues, Stud. Health Technol. Inform. 96 (2003) 15–20.
[33] Y.Y. Han, J.A. Carcillo, S.T. Venkataraman, R.S. Clark, R.S. Watson, T.C. Nguyen, H. Bayir, R.A. Orr, Unexpected increased mortality after implementation of a commercially sold computerized physician order entry system, Pediatrics 116 (2005) 1506–1512.
[34] Standard Guide for Information Access Privileges to Health Information, ASTM E1986–98, 2005.
[35] ISO DTS 21298 Functional and Structural roles, ISO.
[36] E. Coiera, R. Clarke, e-Consent: the design and implementation of consumer consent mechanisms in an electronic environment, J. Am. Med. Inform. Assoc. 11 (2004) 129–140.
[37] T. Adams, M. Budden, C. Hoare, H. Sanderson, Lessons from the central Hampshire electronic health record pilot project: issues of data protection and consent, BMJ 328 (2004) 871–874.
[38] J. Hewison, A. Haines, Overcoming barriers to recruitment in health research, BMJ 333 (2006) 300–302.
[39] O.F. Norheim, Soft paternalism and the ethics of shared electronic patient records, BMJ 333 (2006) 2–3.
[40] P.R. Cundy, A. Hassey, To opt in or opt out of electronic patient records? Isle of Wight and Scottish projects are not opt out schemes, BMJ 333 (2006) 146.
[41] N. Watson, Patients should have to opt out of national electronic care records: FOR, BMJ 333 (2006) 39–40.
[42] J. Wilkinson, Patients should have to opt out of national electronic care records: what's all the fuss about? BMJ 333 (2006) 42–43.
[43] C. Junghans, G. Feder, H. Hemingway, A. Timmis, M. Jones, Recruiting patients to medical research: double blind randomised trial of "opt-in" versus "opt-out" strategies, BMJ 331 (2005) 940.
[44] BMA statement on Connecting for Health (last accessed 2008/02/18); Available from: http://www.bma.org.uk/ap.nsf/Content/cfhstatement.
[45] Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data (last accessed 2008/02/18); Available from: http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:HTML.
[46] M. van der Haak, A.C. Wolff, R. Brandner, P. Drings, M. Wannenmacher, T. Wetter, Data security and protection in cross-institutional electronic patient records, Int. J. Med. Inform. 70 (2003) 117–130.
[47] Good Medical Practice (last accessed 2008/02/18); Available from: http://www.gmc-uk.org/guidance/good_medical_practice/index.asp.
[48] Nationaal ICT Instituut in de Zorg (last accessed 2008/02/18); Available from: http://www.nictiz.nl/.
[49] S. Nouwt, Beveiliging van het EPD, in Rapportage van het juridisch laboratorium, ZonMW/ICZ, Den Haag, p. 65.
[50] B. Blobel, Advanced and secure architectural EHR approaches, Int. J. Med. Inform. 75 (2006) 185–190.
[51] Implementatiehandleiding HL7v3 Zorg Informatie Makelaar (last accessed 2008/02/18); Available from: http://www.nictiz.nl/uploaded/FILES/AORTA%20release%20augustus%202006/hl7_zim%20v2.5.pdf.

[52] RFC 3881 Security Audit and Access Accountability Message XML Data Definitions for Healthcare Applications (last accessed 2008/02/18); Available from: http://tools.ietf.org/html/rfc3881.
[53] Machbarkeitsstudie betreffend Einfuhrung der elektronischen Gesundheitsakte (ELGA) im osterreichischen Gesundheitswesen (last accessed 2008/02/18); Available from: http://www.arge-elga.at/fileadmin/user_upload/uploads/download_Papers/Arge_Papers/Machbarkeitsstudie_ELGA_Endbericht_21112006.pdf.
[54] Architectuurontwerp basisinfrastructuur in de zorg, versie 4.2 (last accessed 2008/02/18); Available from: http://www.nictiz.nl/uploaded/FILES/AORTA%20release%20augustus%202006/Architectuur%20AORTA%20versie%204.2.pdf.
[55] W.W. Simons, K.D. Mandl, I.S. Kohane, The PING personally controlled electronic medical record system: technical architecture, J. Am. Med. Inform. Assoc. 12 (2005) 47–54.
[56] HL7 Message Development Framework (last accessed 2008/02/18); Available from: http://www.hl7.org/library/mdf99/mdf99.pdf.
[57] Peer-To-Peer Networking (last accessed 2008/02/15); Available from: http://en.wikipedia.org/wiki/Peer-to-peer.