

# An Information Accountability Framework for Shared eHealth Policies

Randike Gajanayake  
Queensland University of Technology  
Brisbane, Australia

g.gajanayake@qut.edu.au

Renato Iannella  
NEHTA  
Brisbane, Australia

renato.iannella@nehta.gov.au

Tony Sahama  
Queensland University of Technology  
Brisbane, Australia

t.sahama@qut.edu.au

## ABSTRACT

Privacy issues have hindered the evolution of e-health since its emergence. Patients demand better solutions for the protection of private information. Health professionals demand open access to patient health records. Existing e-health systems find it difficult to fulfill these competing requirements. In this paper, we present an information accountability framework (IAF) for e-health systems. The IAF is intended to address privacy issues and their competing concerns related to e-health. Capabilities of the IAF adhere to information accountability principles and e-health requirements. Policy representation and policy reasoning are key capabilities introduced in the IAF. We investigate how these capabilities are feasible using Semantic Web technologies. We discuss with the use of a case scenario, how we can represent the different types of policies in the IAF using the Open Digital Rights Language (ODRL).

## Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: Security and Protection - *access control*; D.3.m [Programming Languages]: Miscellaneous; K.4.1 [Public Policy Issues]: Privacy; K.5.1 [LEGAL ASPECTS OF COMPUTING]: Hardware/Software Protection - *proprietary rights*.

## General Terms

Management, Design, Security, Human Factors, Standardization, Languages

## Keywords

E-health, Semantic Web, ODRL, Privacy, Information Accountability

## 1. INTRODUCTION

E-health is the use of Information and communications technology (ICT) in healthcare. Amongst others, the Internet is the primary mode of communication for e-health applications. The Web is gradually transforming to what is called “the Semantic Web” where the traditional Syntactic Web is leveraged towards a distributed knowledge repository. The semantic web is based on

the Resource Description Framework (RDF) [1] for metadata semantics and the Web Ontology Language (OWL) [2] for web ontologies. These technologies enable the development of Web based information systems that are capable of automated reasoning, impossible with the syntactic web. These capabilities open new avenues for e-health systems. But, with the use of the Internet to manage health information, the existing concerns in healthcare such as information security and informational privacy become paramount issues needing rigorous attention. This raises questions as to what the relevant security measures are and how an assurance of privacy can be given to the stakeholders (patients and healthcare professionals). In this paper we present an information accountability framework (IAF) for e-health systems. This framework will make applications such as the one proposed by Gajanayake et al. [3] practicable. We consider requirements of different stakeholders in healthcare and accordingly construct our IAF adhering to information accountability principles in the healthcare context.

The rest of this paper is organised as follows. In the next section we will discuss privacy and its impact on e-health. In section 3 we give a brief account on information accountability and the principles behind the concept. In section 5, we present an IAF for e-health systems by extending an access control model from recent work which is summarised in section 4. Section 6 discusses how the introduced capabilities are attainable with available technologies. We will use a simple case scenario to operationalise the concept.

## 2. E-HEALTH AND PRIVACY

An eHR is a complete record of a patient’s medical history. They may also include information pertaining to sensitive concerns such as sexual health, mental health, addictions to drugs or alcohol, abortions etc. Hence unlawful disclosure of personal information could cause the subject of the information embarrassment and may affect insurability, child custody cases, and even employment [4, 5]. Therefore, informational privacy is vital to ensure the reliability of eHR systems. As a result patients demand strong security for their eHRs. Definitions for privacy come in many different forms. Alan Westin, in his book “*Privacy and Freedom*”, defines privacy as “*the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others*” [6], i.e. control of private information. A considerable degree of control over one’s personal information is an essential aspect to protecting information privacy [7]. Due to the disparity of data ownership in healthcare, giving control of the data must be handled with care.

Various methods have been proposed to address the privacy conundrum ranging from strict access control to privacy-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DUMW2012, April 16 2012, Lyon, France.

Copyright 2012 ACM 1-58113-000-0/00/0010...\$10.00.

preserving algorithms. Access control mechanisms either permit or deny access, there are no intermediate states. They are not policy-aware and may also hinder the actions of legitimate users of an information system [8]. According to Kagal et al. [9] relying solely on access control mechanisms to guard information would be inadequate for privacy protection.

Information accountability (IA) can complement access control mechanisms and support policy-awareness. The principles behind IA, in theory, would make sure that information users follow the appropriate rules and policies. To facilitate IA principles, systems should implement usage policies on its assets. Considering data in eHRs digital assets digital rights management (DRM) techniques can be used for the management of the data. Privacy policies in e-health can be represented using an appropriate digital rights expression language (REL). Policies on the use of data in an eHR can be set by the patient, a trusted healthcare representative, a health authority or all the above.

### 3. INFORMATION ACCOUNTABILITY

Accountability systems lack formal foundations making it an attractive theme to many [10-13]. Jagadeesan et al. [10] assume that the relevant privacy policies exist and develop formal foundations for information accountability in terms of the privacy policies which define appropriate sharing of information among agents and provide algorithms that can be used by an auditor to check for compliance with rules. Weitzner et al. [14] propose a solution to the question of compliance of privacy policies by tracking all transactions and making them transparent hence creating an incentive for the users to abide by the rules. They assume that appropriate policy rules exist with a formal representation, policy-aware transaction logs and a policy-reasoning capability which would enable accountability systems to hold information users accountable for misuse. Focusing on the facts Weitzner et al. [14] put forth, Sloan et al. [13] address information accountability in terms of both social policies and technical aspects. They point out difficulties to developing accountability systems by stating that automated checking for compliance of privacy policy is a necessity for accountability systems and without the adequate foundations in both formal models and public policy issues they are unlikely to do so. They believe that policies required to developing accountability systems are informational norms and state that a proper balance between privacy requirements and competing concerns is necessary to sustain the architectural and social aspects introduced by Weitzner et al. [14].

Access control and accountability are closely related concepts. Access control is about restrictions, whereas accountability is about punishment. Hence for accountability systems, *audit logs* are essential [15]. Accountability systems facilitate *fair use* of information. Rather than prevention via rigid locks on data, accountability is about *deterrence*. The presence of an accountability mechanism delivers a threat of punishment which would deter users from intentional misuse. Accountability systems should facilitate *transparency* such that all relevant parties have the capability to observe how information is used and by whom. This makes *bad acts visible* and helps deter users from misuse [14]. The users of an accountability system should be *well informed*, i.e. a notification process where users are informed about underlying policies before an action occurs should be in place. For example a user will be notified whether he is actually authorised to use a particular set of data he is trying to use and the

ramifications if he proceeds regardless of the usage policies in place. This will also help in facilitating non-repudiation which is a significant aspect in information security. When holding someone accountable, the trustworthiness of the data about the inappropriate transaction is critical. Hence, *provenance* of data and metadata is a significant factor in information accountability. Electronic data does not have the necessary historical information that would help end-users, reviewers or regulators make the necessary verifications [16]. In an accountability system provenance can be facilitated using appropriate transaction logs. These transaction logs also serve another purpose in terms of accountability by being *policy-aware*. Policy-aware transaction logs can also facilitate *policy reasoning* capabilities and enable the users to reason about misuse and against claims of misuse.

Creating proper *incentives* that would make consumers follow rules of accountability systems is important [13]. For an information user, the threat of punishment is an incentive to follow system rules. An incentive such as a strong assurance of privacy should be given to patients to prevent them from withholding information or enforcing rigid restrictions on data.

#### 3.1 Information accountability in healthcare

In order to understand the concept of information accountability in healthcare, it is important to clearly identify the different parties in healthcare that can be held accountable, the issues for which a party can be held accountable and the appropriate mechanisms for accountability in healthcare [17]. Policies should be developed that address the different capabilities of roles within the industry. These policies should capture the requirements of all relevant parties. As stated above, in the healthcare domain it is difficult to define who owns health information. It is clear that patients are the subjects of health information. Patients are not always medical professionals; hence it is impossible to give them full control of their health information. Privacy policies should accompany an input from a professional health body such as a trusted medical practitioner or a central health authority. But is it important to balance between the patient's privacy requirements and the requirements of the healthcare providers or the care givers (competing concerns). In a healthcare setting the patients privacy policies cannot contradict those set by the healthcare providers or the health authority. The IMIA code of ethics for medical information professionals [18] states under their first ethics principle; *Principle of Information-Privacy and Disposition* that "*all persons have a fundamental right to privacy, and hence to control over the collection, storage, access, use, communication, manipulation and disposition of data about themselves*". A patient with an eHR, hence, should have the following capabilities; 1) the capability to allow a selected group of medical professionals to access the eHR, 2) the capability to hide certain health information from particular health practitioners who already have access to their eHR, 3) the capability to view and how the data in the eHR is used by authorised personnel, 4) the capability to inquire about potential misuse of data. The data consumers (health professionals and health authority) also have particular requirements. We can identify them as follows; 1) the capability to define security policies within the organization, 2) access to the relevant information in a non-restrictive and timely manner, 3) the capability to share patient health information with other health specialists, 4) the capability to override patients' security settings in special circumstances (e.g. life threatening emergency situations, mental health related situations). It is important to note that usage policy enforcement might not always

be beneficial to the patient. While fulfilling these privacy requirements under no circumstance must the health of the patient be compromised. A clear procedure for overriding usage policies in emergency situations should be defined. The nature of the healthcare domain may force the implementation of a *break the glass* approach in emergency situations. The policy formulation process must consider the requirements of both parties. A compromise between these requirements must entail the final policy representation of the systems and the proper integration of these policies would improve patient confidence in the system.

Apart from the requirements stated above, certain circumstances might require some health conditions be kept hidden from the patients. For example this may be the case for patients suffering from severe mental health conditions where the knowledge of particular illnesses may aggravate existing health conditions. They may also be considered unfit to manage their eHR. We acknowledge this eventuality but consider them as rare occurrences and do not integrate such capabilities in to the framework. However, in such cases the control over the patient's eHR may be given to a custodian or a trusted health professional (HP) such as the patient's GP who can take the patient's role in controlling the eHR.

#### 4. PRIVACY ORIENTED ACCESS CONTROL FOR EHR

Following is a brief description of the access control model in [19]. The access control model uses a combination of the principles behind MAC, DAC, RBAC and PBAC. This was done in order to fulfill the requirements of different stakeholders discussed above. The basic protocol for the proposed access control system is illustrated in Figure 1. We assume that the patient has a comprehensive eHR under a relevant health authority.

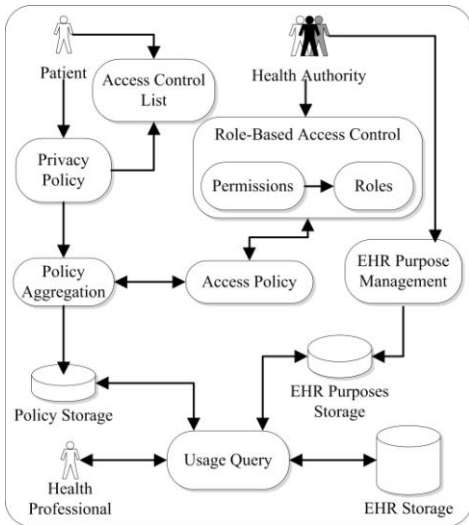


Figure 1. Privacy oriented access control model architecture

The eHR is formulated such that each type of data in the eHR (e.g. identity data, general health data, dental health data, mental health data, etc.) can be distinguished by eHR data type identifiers. For each of these data types there exists a set of predefined purposes for which the data can be used that are defined by a central health authority. The patient and the health

authority can set privacy and access policies respectively. These two policies are later combined using the *Policy Aggregation* to form the final operational policy. The protocol for the policy formulation is as follows.

The health authority defines intended purposes and sensitivity labels for each data type and element. We use object sensitivity labeling using a tree structure (a sensitivity tree (ST)) that has the eHR itself as the root element, the data types as children and data elements as grandchildren. A sensitivity label is not assigned to the objects themselves rather we relate the access level of a particular user (health professional) in terms of the sensitivity label of the data elements. Note that the sensitivity labels mentioned here are different from the classical hierarchical security levels found in MAC [20]. The nature of health information makes it difficult to define a clear hierarchical structure for the sensitivity of data elements that is general to all patients. For example, sexual health and mental health information may have the same sensitivity for some patients and may not be so for others.

**Definition:** A sensitivity label (SL) is a tuple  $\langle \text{ASL}, \text{PSL} \rangle$ , where  $\text{ASL} = \{\text{asl}_1, \text{asl}_2, \dots, \text{asl}_n\}$  is a set of allowed sensitivity labels and  $\text{PSL} = \{\text{psl}_1, \text{psl}_2, \dots, \text{psl}_n\}$  is a set of prohibited sensitivity labels.

$\text{ASL} = \{\text{asl}_i\}; i = 1 \dots n$  is denoted as all of the descendants of  $\text{asl}_i$  including  $\text{asl}_i$ .

$\text{PSL} = \{\text{psl}_j\}; j = 1 \dots n$  is denoted as all of the descendants of  $\text{psl}_j$  including  $\text{psl}_j$ .

**Example:** Matt can access Gary's mental health details but cannot access his Sexual or Dermatology details. The access level for Matt can be represented in terms of sensitivity labels as follows.

$$\text{SL}_{\text{Matt}} = \langle \{\text{eHR}\}, \{\text{Sexual Health, Dermatology Health}\} \rangle$$

Here we use the Denial-Takes-Precedence principle. Access is granted to the entire eHR and then access is denied to specific field by the PSL. This helps isolate the most sensitive information in the eHR that need to be hidden from certain users.

Patients and the health authority set sensitivity levels for health professionals. The sensitivity level defined by the health authority is different to the ones defines by the patients. PSLs set by the health authority will always be *NULL*. This is because the health authority is concerned with allowing access to data elements. The prohibitions are defined by the patients. The ASL set by the patients always precedes that which is set by the health authority. The ASL set by the health authority always precedes PSL set by the patients if there is a conflict. This feature will ensure that the relevant information is always available to the relevant health professional.

#### 5. INFORMATION ACCOUNTABILITY FRAMEWORK FOR E-HEALTH

Here we present an information accountability framework (IAF) for e-health systems. It can be considered as an extension to the access control model described above. In the IAF the policies defined in the access control model act as the underlying policies to which the users must comply to but do not prevent users from accessing data. This is to facilitate unrestricted access to health information for authorised users. The reasoning capability of the IAF takes these policies in to consideration whilst performing such tasks.

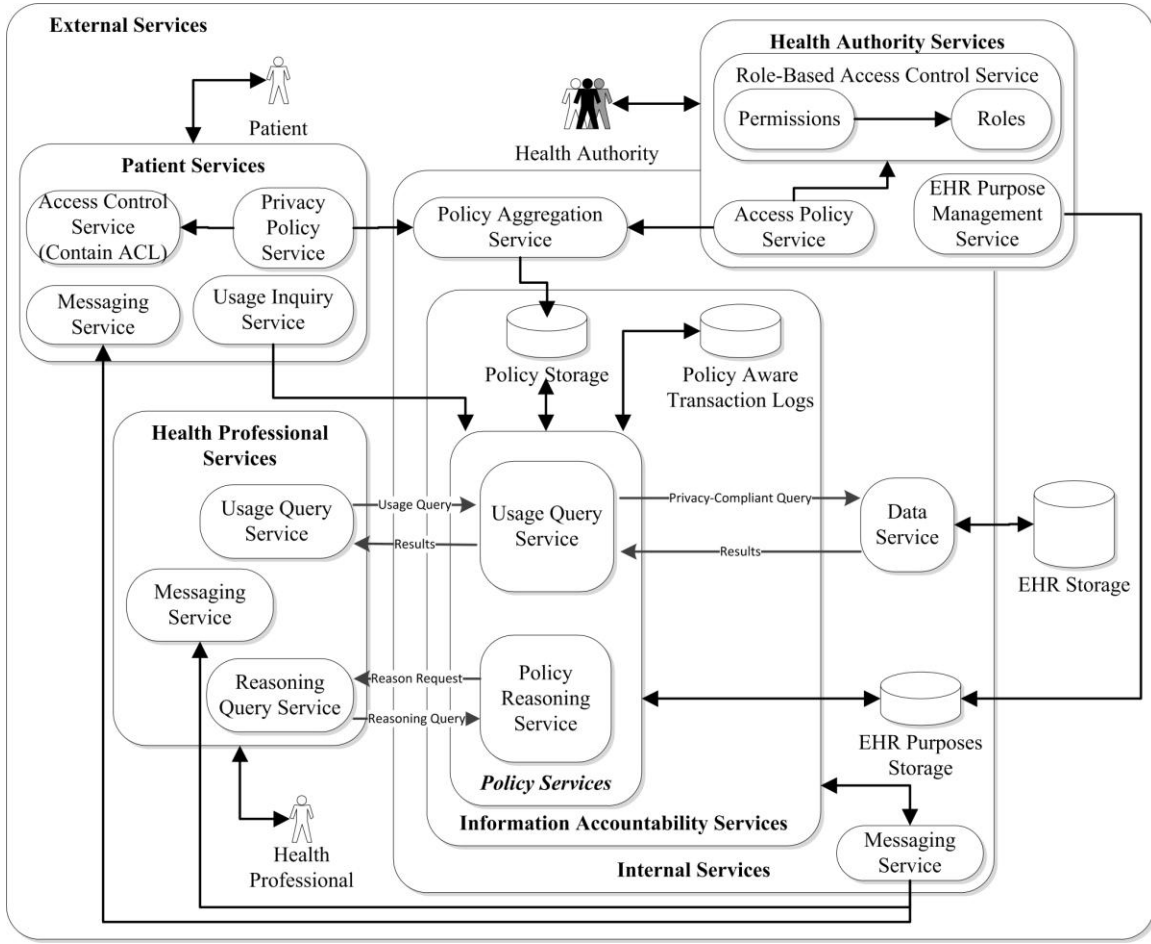


Figure 2: Schematic IAF architecture

The IAF is divided into two categories of services; external services and internal services and have three types of users; patients ( $P$ ), a health authority ( $HA$ ) and health professionals ( $HP$ ). A schematic architecture is shown in the Figure 2.

Internal services consist of a *policy aggregation service*, the *information accountability services*, a *messaging service*, a *data service*, *policy storage* and the *EHR Purpose storage*. External services of the IAF include *patient services*, *health authority services*, *health professional services* and the external *EHR storage*. Detailed descriptions of these services are given next.

### 5.1 Internal services

Information accountability services consists of *policy storage* ( $PS_{IAS}$ ), *policy aware transaction logs* ( $PATL_{IAS}$ ) and policy services containing a *usage query service* ( $UQS_{IAS}$ ) and a *policy reasoning service* ( $PRS_{IAS}$ ).  $PS_{IAS}$  stores the policies it receives from the *policy aggregator service*.  $UQS_{IAS}$  processes the usage queries it receives from health professional services requesting access to EHR data. Once the policy service receives an *inquiry query* from patient services  $PRS_{IAS}$  send a request to the health professional service requesting a *reasoning query* for a particular information usage instance. The reasoning queries are processed with the use of  $PATL_{IAS}$  which contains all past transactions of the system.

Other internal services include a *policy aggregator service* ( $PAS_{IS}$ ) which amalgamates the policies from  $PPS_p$  and  $APS_{HA}$  in

such a way that the patient's privacy requirements are met and the health authorities' policies be satisfied, a *data service* ( $DS_{IS}$ ) which is the only component with access to the EHR storage, a *messaging service* ( $MS_{IS}$ ) which sends out the relevant messages to other services and an *EHR purposes storage* ( $EPS_{IS}$ ) which consists of the intended purposes of each of the data types in the EHR. The  $EPS_{IS}$  is managed by  $HA$ .

### 5.2 External services

External services are used by the end users to give inputs to the internal services and receive results from them. External services consist of patient services, health authority services, health professional services and the EHR storage.

Patient services are used by a patient to manage their EHR. The patient services consist of an *access control service* ( $ACS_p$ ), *privacy policy service* ( $PPS_p$ ), *messaging service* ( $MS_p$ ) and a *usage inquiry service* ( $UIS_p$ ). A patient maintains an access control list (ACL) with the use of  $ACS_p$ . The patients set their privacy policies using  $PPS$  and assign sensitivity levels for trusted health professionals in the ACL. These policies are then amalgamated by the *policy aggregation service* ( $PAS_{IS}$ ) with the policies of the health authority and stored in  $PS_{IAS}$ . Patients receive notifications and can send messages to HPs through the  $MS_p$  from the internal services. Notifications include regular updates on the EHR, notifications of information access by HPs, warnings of potential information misuse and messages from HPs.

All messages need to go through the internal services for them to be recorded in the Transaction logs.

Health authority services are used by a central health authority to manage access settings for health professionals. Health authority services consists of a *role based access control service* ( $RBACS_{HA}$ ), an *EHR purpose management service* ( $EPMS_{HA}$ ) and *access policy service* ( $APS_{HA}$ ). The HA set minimum access levels for HPs using  $APS_{HA}$  together with  $RBACS_{HA}$ . These policies are combined with the patient's privacy policies according to the access control protocol in [19], which is also summarised in section 3. HA uses  $EPMS_{HA}$  to manage the EHR purposes in  $EPS_{IS}$ .

Health professional services are used by health professionals to access patient EHR information. HPs are able to perform actions such as read and write. HPs are also able to initiate information sharing requests in order to share patient health information with other HPs to make informed decisions. Health professional services include a *usage query service* ( $UQS_{HP}$ ), a *reasoning query service* ( $RQS_{HP}$ ) and a *messaging service* ( $MS_{HP}$ ). HPs can lodge usage queries using  $UQS_{HP}$  requesting access to EHR information. These queries contain purposes for which information is required. The queries are processed by the  $UQS_{IAS}$  and if they are policy compliant access is granted. If the usage queries are not policy compliant a warning notification is sent to the requester at which point he can either comply with the warning or disregard it. If the warning is disregarded and the data is accessed by the HP, a message is sent by the  $MS_{IS}$  to  $MS_P$  notifying the patient of potential information misuse. At this point the patient may initiate a usage inquiry using  $UIS_P$ . As a result  $PRS_{IAS}$  sends a request to  $RQS_{HP}$ . The HP then has to send a justification of the use of information in the form of a reasoning query through the  $RQS_{HP}$ . The justification is processed by the  $PRS_{IAS}$ . If the provided justification is valid the incident is resolved. If not, further action (such as legal action) would be taken which we would not discuss in this paper (a justifiable action would be in the case of an emergency where the existing policies had to be overridden for the sake of the patient's health).  $PRS_{IAS}$  should have the capability to deduce whether a provided justification is valid. This process of inquiries and resulting justifications enables the system to detect intentional misuse of data by users.

## 6. FEASIBILITY

The main capabilities of the IAF are policy representation, policy storage and policy reasoning capabilities. The key challenge in implementing the IAF in a technical point of view is to fulfill these capabilities. In this section we will give an account as to how these capabilities are feasible through available semantic web technologies.

As discussed in section 3, proper representation of policies is vital in information accountability. For our model we look to digital rights management (DRM) as a solution. Apart from their applications in copyright protection of media files, etc on the Internet, DRM technologies are becoming a prominent resource in protecting private information of individuals [21]. DRM has many similarities to the traditional access control model but differs in that they require information to remain protected even after access is granted to authorised users. DRM deals with usage control of a piece of information resource by authorised users. Each piece of information is protected by a usage license created by the digital rights holder. DRM can benefit e-health technologies by

providing a means to manage the use of eHRs. Rights expression languages (REL) are a critical aspect of DRM systems. The Open Digital Rights Language (ODRL) version 2 [22] is based on XML and provides a syntax and semantics to express policies related to digital assets. We have chosen ODRL as the policy language for our model because it is independent of implementation constraints and it's capable of expressing a wide range of policy-based information.

### 6.1 Healthcare scenario

Consider the following scenario. Gary has a comprehensive eHR. Gary has a list of trusted healthcare providers (health professionals) to whom he gives access to data in his eHR. Peter is Gary's GP, Sandra is a dermatologist, Bill is a sexual health specialist and Matt is a mental health specialist who has treated Gary in the recent past. Gary can set privacy settings to govern the access to his eHR. A central health authority can also set access settings to patient's eHR by considering the roles of each health professional. In addition to privacy and access policies, other constraints can be present in the eHR. One such policy can be a *take control* policy. In which an eHR holder may take control of their eHR at the age of 15 (which was previously controlled by a parent, legal guardian of authorised representative) and must take control at the age of 18. Such policies must accompany privacy and access policies in the eHR.

#### 6.1.1 Scenario

After noticing a skin rash, Gary visits his trusted dermatologist Sandra for a check up. The preliminary examination reveals that Gary's skin condition could be linked to a known sexually transmitted disease (STD). Gary does not have a sexual health specialist in his list of trusted health professionals. However, Sandra wants to share Gary's details with a sexual health specialist, Bill, in order to get a specialists opinion on the situation. Bill has a default access level set by the health authority to be able to access patients' sexual health details and dermatology details. Since Sandra is in Gary's list of trusted HPs to be able to access Gary's dermatology information, she can initiate a request to share Gary's details with other health professionals. Gary, however, is notified of this action by Sandra. After Bill gets this request, he initiates a usage request to use the data for diagnosis purposes. At some point during or after this episode of care, Gary may include Bill to his list of trusted health professional.

#### 6.1.2 ODRL policies

Gary allows Sandra to access his EHR but restricts her from accessing his sexual health details and mental health details. Below is an ODRL V2 XML instance of this policy.

```
<o:policy xmlns:o="http://w3.org/ns/odrl/2"
xmlns:eh="urn:health.gov" type="
http://w3.org/ns/odrl/2/privacy" uid="policy-use-ehr"
conflict="o:prohibit">
  <o:permission>
    <o:asset uid="urn:ehr:gary" relation="o:target"/>
    <o:party uid="urn:patient:gary" role="o:assigner"/>
    <o:party uid="urn:healthPro:dermatHealth:sandra"
      role="o:assignee"/>
    <o:action name="o:read"/>
    <o:constraint name="o:purpose" operator="o:eq"
      rightOperand="eh:healthCare"/>
  </o:permission>
```

```

<o:prohibition>
  <o:asset uid="urn:chr:gary" relation="o:target"/>
  <o:party uid="urn:patient:gary" role="o:assigner"/>
  <o:party uid="urn:healthPro:dermatHealth:sandra"
    role="o:assignee"/>
  <o:action name="o:read"/>
  <o:constraint name="o:purpose" operator="o:eq"
    rightOperand="ch:sexualHealthCare"/>
  <o:constraint name="o:purpose" operator="o:eq"
    rightOperand="ch:mentalHealthCare"/>
</o:prohibition>
</o:policy>

```

The conflict attribute of the policy above is set to “*prohibit*” indicating that prohibitions take precedence in the policy. The health authority can set an access policy for Sandra which is given below.

```

<o:policy xmlns:o="http://w3.org/ns/odrl/2"
xmlns:ch="urn:chealth.gov" type="
http://w3.org/ns/odrl/2/agreement" uid="policy-use-ehr">
  <o:permission>
    <o:asset uid="urn:chr:gary" relation="o:target"/>
    <o:party uid="urn:health:authority" role="o:assigner"/>
    <o:party uid="urn:healthPro:dermatHealth:sandra"
      role="o:assignee"/>
    <o:action name="o:read"/>
    <o:constraint name="o:purpose" operator="o:eq"
      rightOperand="ch:dermatHealthCare">
    <o:constraint name="o:purpose" operator="o:eq"
      rightOperand="ch:sexualHealthCare">
  </o:permission>
</o:policy>

```

The health authority is responsible for setting default access policies for healthcare roles, in this case for the role of a dermatologist. In the policy above HA gives Sandra the permission to access Gary’s dermatology details and sexual health details. Note here that Gary’s settings prohibit Sandra from accessing his sexual health details. But we assume a hypothetical scenario where a relationship between skin conditions and STDs exist, and every dermatologist should have access to the patient’s sexual health details. The health authority is aware of this fact and allows all dermatologists access to patients sexual health details. According to the access control protocol in section 3, the settings by the health authority always prevail over patient settings. The final policy will be a combination of the two policies and hence the requirement for  $PAS_{IS}$  in the IAF. The amalgamated policy for Sandra is given below.

```

<o:policy xmlns:o="http://w3.org/ns/odrl/2"
xmlns:ch="urn:chealth.gov" type="
http://w3.org/ns/odrl/2/privacy" uid="policy-use-ehr"
conflict="o:prohibit">
  <o:permission>
    <o:asset uid="urn:chr:gary" relation="o:target"/>
    <o:party uid="urn:patient:gary" role="o:assigner"/>
    <o:party uid="urn:healthPro:dermatHealth:sandra"
      role="o:assignee"/>
    <o:action name="o:read"/>
    <o:constraint name="o:purpose" operator="o:eq"
      rightOperand="ch:healthCare"/>
  </o:permission>
  <o:prohibition>
    <o:asset uid="urn:chr:gary" relation="o:target"/>
    <o:party uid="urn:patient:gary" role="o:assigner"/>

```

```

    <o:party uid="urn:healthPro:dermatHealth:sandra"
      role="o:assignee"/>
    <o:action name="o:read"/>
    <o:constraint name="o:purpose" operator="o:eq"
      rightOperand="ch:mentalHealthCare"/>
  </o:prohibition>
</o:policy>

```

This final policy is stored in  $PS_{IAS}$  and is used by other services. Updates are done to the policies in  $PS_{IAS}$  accordingly.

Information sharing is an important aspect of healthcare and is facilitated in the IAF. HPs who are already in the ACL can initiate sharing requests.

```

<o:policy xmlns:o="http://odrl.net/2.0"
xmlns:ch="urn:chealth.gov" type="
http://w3.org/ns/odrl/2/request" uid="policy-share-ehr">
  <o:permission>
    <o:asset uid="urn:chr:gary" relation="o:target"/>
    <o:party uid="urn:healthPro:dermatHealth:sandra"
      role="o:assignee"/>
    <o:action name="o:share"/>
    <o:constraint name="o:purpose" operator="o:eq"
      rightOperand="ch:dermatHealthCare">
    <o:constraint name="o:recipient" operator="o:eq"
      rightOperand="urn:healthPro:dermatHealth:bill">
  </o:permission>
</o:policy>

```

In the policy above Sandra initiates a request to share Gary’s dermatology details with Bill. Bill accepts this request by lodging the following access request to read Gary’s dermatology details. Requests resulting from sharing requests are allowed (holding to general access policies) since the initial request was from a HP already in the ACL.

```

<o:policy xmlns:o="http://w3.org/ns/odrl/2"
xmlns:ch="urn:chealth.gov" type="
http://w3.org/ns/odrl/2/request" uid="policy-use-ehr">
  <o:permission>
    <o:asset uid="urn:chr:gary" relation="o:target"/>
    <o:party uid="urn:healthPro:sexualHealth:bill"
      role="o:assignee"/>
    <o:action name="o:read"/>
    <o:constraint name="o:purpose" operator="o:eq"
      rightOperand="ch:dermatHealthCare">
  </o:permission>
</o:policy>

```

Using ODRL we can formulate the different types of policies within the eHR system. But the Challenge lies in using ODRL in the Semantic Web domain. Next we will look at how we can use ODRL in conjunction with semantic web technologies and how we can attain the capabilities proposed for the IAF.

## 6.2 ODRL in the Semantic Web

ODRL is a solution to move DRM to the Internet. But in order to enforce the semantics of the policies in conjunction with ODRL, a corresponding ontology is required. At the time of writing such ontology was not present. The ontology for the policies can be represented using OWL. Even though a comprehensive ontology for ODRL V2 is required for an end , we will not present one in this paper. Such ontologies allow us to achieve the reasoning capabilities proposed in the IAF.

$EPS_{IS}$  contains an ontology representing the relationships between the eHR data themselves and eHR data and the intended purposes.

This ontology together with a comprehensive medical ontology enables us to infer facts otherwise would not be available. For example, the presence of the fact that Gary has a particular allergy in the  $EPS_{IS}$  can lead to the inference of the fact that a particular medication has the tendency to be harmful to Gary. This fact would not have been available to the eHR system without a specific external input saying so or if Gary has had an illness which is usually treated by this particular medication. The inferences are updated with new data and facts available to  $EPS_{IS}$ . The policies in  $PS_{IAS}$  are stored in RDF with vocabularies from the ODRL ontology. The queries made by  $UIS_p$  and  $PRS_{IAS}$  are made in a RDF query language like SPARQL [23]. Data stored in  $PATL_{IAS}$  is also in RDF allowing mining to be done using SPARQL. Together with these services and a policy aware reasoner,  $PRS_{IAF}$  allow us (with a suitable natural language translation middleware) to process queries such as “Why did Sandra read my sexual health details?” by Gary. Similarly, Sandra will be able to justify why she read Gary’s sexual health details. The validity of the justification is determined after mining the  $PATL_{IAS}$  and  $PS_{IAS}$ . A provided justification holds if the facts confirm with the available knowledge. Note here that as mentioned above, the patient can only lodge an inquiry query if there has been a possible misuse of data i.e. some underlying policy has been violated by the user. The justification is on why the user has done so. The ontologies defined enable us to infer facts that validate the justification. For example, in an emergency situation the treating health professional will access all necessary information from the eHR regardless of the privacy and access policies. This will be recorded in  $PATL_{IAS}$ . For any inquiry made by the patient to clarify data usage related to this episode of care, the fact that the incident was considered and recorded as an emergency would validate the justifications given by the health professionals.

## 7. CONCLUSION AND FUTURE WORK

We have presented an information accountability framework (IAF) for e-health which adheres to information accountability (IA) principles and requirements of stakeholders in healthcare. IA is a term better defined contextually rather than in a general sense. We focused on the healthcare context and treated each IA principle accordingly. The requirements of the healthcare domain which we considered are mainly privacy requirements of patients and access and usage requirements of health professionals. Amongst others these carry the most potential to hinder the development of e-health systems and are the main concerns of consumers of those systems. In any accountability system, policy representation is clearly a key aspect. In our model we used ODRL as the policy language and discussed how we can represent the different privacy and access policies in the IAF. Semantic Web based policy management has been studied by many and some attractive solutions have been proposed [24-26]. However, we chose ODRL as the policy language for our model to give us the flexibility needed to extend the existing model to suit the capabilities introduced in our model. For example using an extension to ODRL we can represent policy aware transaction logs which are then used for reasoning to detect misuse.

Policy reasoning is the other key factor in information accountability. Currently the only technologies that provide such capabilities and are readily available are semantic web technologies. We discussed how we can use semantic web technologies such as OWL ontologies and RDF to develop the proposed IAF. It is clear that developing a comprehensive eHR

system with an IAF is an immense undertaking. But with the level of technology currently at the disposal of developers it is without a doubt feasible task.

In e-health, accountability systems will enable the use of health information in a more free but controlled manner. This will allow health professionals to access relevant information at any point without the restrictions currently present in e-health solutions. We believe that the presence of the IAF will increase the confidence of the patients towards e-health systems and would lead to e-health systems being better adopted. Barriers still exist in our venture towards building a working system with the capabilities introduced. We are currently working on demonstrating the presented IAF using the technologies discussed. At the time of writing the development of the aforementioned ontologies are ongoing. Building a comprehensive eHR system is not our goal. Our goal is to show that with IA capabilities the current state of e-health systems can be improved to a more open and healthcare oriented state from a security and privacy oriented state.

## 8. ACKNOWLEDGEMENTS

We would like to thank the National Information and Communications Technology Australia (NICTA) for partially funding this ongoing research project.

## 9. REFERENCES

- [1] Lassila, O. and Swick, R. *Resource Description Framework (RDF) Model and Syntax Specification*. 1999.
- [2] McGuinness, D. and van Harmelen, F. *OWL Web Ontology Language Overview*. City, 2004.
- [3] Gajanayake, R., Iannella, R. and Sahama, T. Sharing with Care: An Information Accountability Perspective. *Internet Computing, IEEE*, 15, 4 (2011), 31-38.
- [4] Pratt, W., Unruh, K., Civan, A. and Skeels, M. M. Personal health information management. *Commun. ACM*, 49, 1 (2006), 51-55.
- [5] Cannoy, S. D. and Salam, A. F. A framework for health care information assurance policy and compliance. *Commun. ACM*, 53, 3 (2010), 126-131.
- [6] Westin, A. *Privacy and Freedom*. New York Atheneum, 1967.
- [7] Solove, D. J. Understanding Privacy. *Daniel J. Solove, UNDERSTANDING PRIVACY, Harvard University Press, May 2008*(2008).
- [8] Kagal, L. and Pato, J. Preserving Privacy Based on Semantic Policy Tools. *Security & Privacy, IEEE*, 8, 4 (2010), 25-30.
- [9] Kagal, L. and Abelson, H. *Access Control is an Inadequate Framework for Privacy Protection*. City, 2010.
- [10] Jagadeesan, R., Jeffrey, A., Pitcher, C. and Riely, J. *Towards a Theory of Accountability and Audit Computer Security – ESORICS 2009*. Springer Berlin / Heidelberg, City, 2009.
- [11] Feigenbaum, J., Hendler, J., Jaggard, A. D., Weitzner, D. J. and Wright, R. N. *Accountability and Deterrence in Online Life*. City, 2011.
- [12] Feigenbaum, J., Jaggard, A. D. and Wright, R. Towards a Formal Model of Accountability. In *Proceedings of the New Security Paradigms Workshop* (CA, USA, September 12-15, 2011), [insert City of Publication],[insert 2011 of Publication].

- [13] Sloan, R. H. and Warner, R. Developing Foundations for Accountability Systems: Informational Norms and Context-Sensitive Judgments. *Annual Computer Security Applications Conference, Workshop on Governance of Technology, Information, and Policies*, 2010(2010).
- [14] Weitzner, D. J., Abelson, H., Berners-Lee, T., Feigenbaum, J., Hendler, J. and Sussman, G. J. Information accountability. *Commun. ACM*, 51, 6 (2008), 82-87.
- [15] Lampson, B. Privacy and security: Usable security: how to get it. *Commun. ACM*, 52, 11 (2009), 25-27.
- [16] Moreau, L., Groth, P., Miles, S., Vazquez-Salceda, J., Ibbotson, J., Jiang, S., Munroe, S., Rana, O., Schreiber, A., Tan, V. and Varga, L. The provenance of electronic data. *Commun. ACM*, 51, 4 (2008), 52-58.
- [17] Emanuel, E. J. and Emanuel, L. L. What Is Accountability in Health Care? *Annals of Internal Medicine*, 124, 2 (January 15, 1996 1996), 229-239.
- [18] International Medical Informatics Association *IMIA Code of ethics for health information professionals*. City, 2002.
- [19] Gajanayake, R., Iannella, R. and Sahama, T. Privacy Oriented Access Control for Electronic Health Records. In *Proceedings of the Workshop on Data Usage Management on the Web* (Lyon, France, 2012).
- [20] Sandhu, R. S. and Samarati, P. Access control: principle and practice. *Communications Magazine, IEEE*, 32, 9 (1994), 40-48.
- [21] Feigenbaum, J., Freedman, M. J., Sander, T. and Shostack, A. Privacy Engineering for Digital Rights Management Systems. *Lecture Notes in Computer Science, Security and Privacy in Digital Rights Management*, 2320(2002), 76-105
- [22] ODRL Initiative *ODRL V2.0 - Core Model - Working Draft*. City, 2012.
- [23] Prud'hommeaux, E. and Seaborne, A. *SPARQL Query Language for RDF*. City, 2008.
- [24] Kagal, L., Finin, T. and Joshi, A. *A Policy Based Approach to Security for the Semantic Web*. Springer Berlin / Heidelberg, City, 2003.
- [25] Kagal, L., Finin, T. and Anupam, J. *A policy language for a pervasive computing environment*. City, 2003.
- [26] Tonti, G., Bradshaw, J., Jeffers, R., Montanari, R., Suri, N. and Uszok, A. *Semantic Web Languages for Policy Representation and Reasoning: A Comparison of KAoS, Rei, and Ponder*. Springer Berlin / Heidelberg, City, 2003.