

# Kent Academic Repository

## Full text document (pdf)

### Citation for published version

Lu, Yang and Sinnott, Richard O. (2015) Semantic Security for E-Health: A Case Study in Enhanced Access Control. In: 2015 IEEE 12th Intl Conf on Ubiquitous Intelligence and Computing and 2015 IEEE 12th Intl Conf on Autonomic and Trusted Computing and 2015 IEEE 15th Intl Conf on Scalable Computing and Communications and Its Associated Workshops (UIC-ATC-ScalCom).

### DOI

<https://doi.org/10.1109/UIC-ATC-ScalCom-CBDCCom-IoP.2015.90>

### Link to record in KAR

<https://kar.kent.ac.uk/80963/>

### Document Version

Publisher pdf

#### Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

#### Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

#### Enquiries

For any further enquiries regarding the licence status of this document, please contact:

[researchsupport@kent.ac.uk](mailto:researchsupport@kent.ac.uk)

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

# Semantic Security for e-Health: A Case Study in Enhanced Access Control

Yang Lu & Richard O. Sinnott

Department of Computing and Information System,  
University of Melbourne,  
Melbourne, Australia  
[luy4@student.unimelb.edu.au](mailto:luy4@student.unimelb.edu.au)  
[rsinnott@unimelb.edu.au](mailto:rsinnott@unimelb.edu.au)

**Abstract**— Data collection, access and usage are essential for many forms of collaborative research. E-Health represents one area with much to gain by sharing of data across organisational boundaries. In such contexts, security and access control are essential to protect the often complex, privacy and information governance concerns of associated stakeholders. In this paper we argue that semantic technologies have unique benefits for specification and enforcement of security policies that cross organisation boundaries. We illustrate this through a case study based around the International Niemann-Pick Disease (NPD) Registry ([www.inpdr.org](http://www.inpdr.org)) - which typifies many current e-Health security processes and policies. We show how approaches based upon ontology-based policy specification overcome many of the current security challenges facing the development of such systems and enhance access control by leveraging existing security information associated with clinical collaborators.

**Keywords**— access control; trust management; Niemann-Pick Disease; ontology

## I. INTRODUCTION

With the demands for increased safety and effectiveness of medications, diagnostic products, and disease prevention, collaborative e-health projects have attracted considerable interest and importance over the last decade [1]. Many e-Health projects utilise electronic health record systems to support a range of research demands: public health, patient care through to biomedical/clinical research. To offer secure access to sensitive information, data protection mechanisms are essential [2]. In particular, before allowing (authorising) users to access data/services, trust relationships need to be established. For flexibility and scalability, ideally these trust relationships should be established between *entities* that were potentially unknown to each other before and leveraging a range of pre-existing information to determine access control decisions. These *entities* might be users but more commonly they will be organisations that have agreed to trust one another – to a certain degree! The degree of trust can and indeed often will vary between organisations. The most elementary of trust relationships is that a given organisation accepts the authentication processes adopted and used by a collaborating organisation. Authentication is a typical starting point of security. This federated level of trust is now widely adopted in academia through utilisation of the Internet2 Shibboleth technologies (<https://shibboleth.net/>) to support decentralised authentication models. Authentication only provides basic trust and the clinical domain in particular requires much finer-

grained access control mechanisms – ideally to support policy-based access control (authorisation) models.

In this context, this paper utilises semantic web technologies to support such capabilities. We focus on improved reasoning and establishment of a trust management framework associated with the International Niemann-Pick Disease (NPD) Registry (INPDR - [www.inpdr.org](http://www.inpdr.org)). Niemann-Pick diseases are rare genetic diseases with distinct clinical spectrum and extremely poor prognosis for patients. Currently there is no specific cure for NPD. The scarcity of patients and associated lack of neurological expertise has led to challenges in delayed diagnosis and difficulty in treating patients [5]. The rarity of the condition demands that centralisation (aggregation) of patient information occurs. The INPDR has been established to aggregate such information from patients originating from many centres and countries. The INPDR system has a simple yet powerful security model, however we believe that semantic technology combined with federated authentication could greatly improve the scalability, and ultimately the usability and flexibility of the platform.

The INPDR system has adopted a role-based access control model (RBAC) to restrict access and use of the patient data in the registry. In order to access the registry, users are required to enroll and be assigned a “username/password” and given an associated role. The username is currently based on their existing email address. In this model, both authentication (username/password) and authorisation (the roles they are assigned) are based on a single, centralised model in which INPDR security administrators have to deal with individual requests. Each request requires consultation with existing INPDR collaborators to ratify that the requestor is known and should be given a particular access privilege. Once ratified, a list of pre-defined roles/attributes is used to assign privileges to the new user. This model works, however it has clear limitations when the system scales to deal with hundreds or thousands of clinicians/researchers where they may move from hospitals or be assigned new positions within their organisations. Ideally, compared to this centralised security model, a scalable registry should be able to work with several trusted “attribute authorities” using information (credentials) from distributed participant organisations. Based on this, the registry should be able to grant/deny requestors based on their associated “proofs” of authentication and authorisation (existing privileges they may possess that have been assigned by one or more trusted authorities). Several authors have explored the extension to the Internet2 Shibboleth technology for federated authorisation, however the standardisation of

roles and privileges required for authorisation remains a challenge [6][7].

The semantic web aims at providing an environment in which trust relations can be established between discrete entities through reasoning about the relations between concepts, such as semantic equivalence and subsumption [8]. This is particular applicable to mapping different security vocabularies and in inferring (evaluating) security policies. In this area, ontologies play an essential role in sharing knowledge and specifying policies [9]. This paper proposes an ontology-based authorisation and trust establishment framework, through which agreement between trusted authorities can be achieved and semantic reasoning between different vocabularies undertaken to provide enhanced access control decisions. Underpinning this work is the notion of a degree of trust related to authentication. At present a common platform for federated access control (authentication) in the health domain (or in government or many other domains) does not exist. The technologies to deliver such federated authentication leveraging pre-agreed trust relationships are available however, e.g. Shibboleth and OAuth [10]. We thus propose a model of future e-Health collaborations and not one that is immediately supported and adopted by the clinical research community.

## II. RELATED WORK

Security in distributed systems normally includes two steps: trust establishment and access control. To authenticate entities in a distributed system, public key infrastructures (PKI) are commonly used. Earlier PKIs were designed for authentication and bound a globally unique identity to each public key. In the case of X.509-based PKIs, X.509 public certificates include a public key binding with the distinguished name (DN) of the key holder, to prove the identity of the public key holder. These keys are assigned by a trusted authority (Certificate Authority) that takes steps and processes to ensure that the user (as given by the DN) is who they claim to be. This often builds on local identity establishment steps (through a Registration Authority). Other existing PKIs include privacy enhanced mail (PEM) [11], securing the internet e-mail system, and the Pretty Good Privacy (PGP) [12], providing cryptographic routines for email and file storage. However, trust establishment in the Web is significantly different from that in closed systems, which is mainly based on “identity-capability” mappings. With the expansion of the web and increased number of diverse, remote users and organisations, and associated sources of (more or less) trusted information, the PKI model is being challenged in its scalability and flexibility, or more precisely the PKI model does not tackle finer-grained security demands facing many domains (such as e-Health).

As e-Health systems required data to be accessible/cross organisational boundaries, it is essential that fine-grained security be supported. Many works have been undertaken regarding security in e-Health. To make sure that heterogeneous data access can be given the right “security behaviour”, Rossilawati et al. established a security classification model, in which communications are regulated in different security process contexts [13]. Key to this was avoiding unnecessary authorisation on sensitive resource access. Based on the principle of “least privilege”, workflow based access control (WBAC) has been proposed to refine the

resource instances and adapt access rights bound to the workflow duties [14]. Such policy-driven methods are based on the assumption that messages from service requestors and providers can be understood bilaterally. However, with cross-platform collaborations emerging, the scale of e-Health systems has substantially expanded in the last decade. As such, solutions to this issue including bridging gaps among distributed sites through flexible and heterogeneous security frameworks are essential, i.e. a common security vocabulary cannot be expected to exist. It is also noted that heterogeneous data vocabularies also exist with a variety of coding systems in use across e-Health data providers. A cross-platform secure architecture with multiple national and regional security domains is required to solve such issues [15]. With the help of an “inter-domain zone”, common security and interoperability services such as policy translation, distributed authentication, authorisation as well as auditing services can be established. The platform can use PKIs for authentication, allowing external clinicians (i.e. external to the organisation making data/services available) to access the records shared with the collaborating parties. Bridging technologies can be used to support integration of PKIs between different security (trusted) domains – recognising that this incurs challenges associated with certification path validation [16]. In this paper, we focus on the challenges of re-use and repurposing of authorisation information that exists across collaborating e-Health sites.

Access control requires policy-based authorisations, i.e. “what can be done by users on remote distributed resources”. According to the Trusted Computer System Evaluation Criteria [17], there are two types of access control models: discretionary *access control* (DAC) [18] and *mandatory access control* (MAC) [19]. DAC provides a means of restricting access to objects/resources based on the identity of subjects, leaving the granting and revoking of privileges up to the resource providers. Typically it can be implemented as access control lists (ACLs) and associated Capabilities. MAC is a means of restricting access to objects based on the sensitivity of the information related to the objects and the authorisation of subjects required and enforced as a clearance. RBAC was proposed with the notion of “grouping” of privileges, where users are defined as being members of a group (specifically having a given role) and the privileges are subsequently mapped to it. Different variants of RBAC have been proposed and implemented including Temporal RBAC [20], Location and Time-based RBAC [21] and Spatial RBAC [22]. Although people are attaching richer and richer semantics on these models, they are fundamentally based upon trust and largely static assignment of privileges to individuals. The exact privileges (roles) are required to be used in security policies for local decisions. Security-policies based upon a static set of pre-negotiated and assigned privileges have scalability issues [23].

To illustrate these issues we consider the Internet2 Shibboleth technologies and their support for single sign-on (SSO) across academia, e.g. the Australian Access Federation (AAF - [www.aaf.edu.au](http://www.aaf.edu.au)). The AAF allows academics across Australia to authenticate to a range of distributed services. The model has an underlying PKI that supports trusted communications between sources of identity (Identity Providers) and the services themselves (Service Providers). This model is based upon a pre-negotiated set of *eduPerson*

attributes, which are assigned by organisations within the AAF, e.g. the University of Melbourne. However knowing that a researcher is an academic from the University of Melbourne is not enough to make a decision to allow access to a resource such as the INPDR. Whilst the authentication may be trusted, the finer-grained roles and privileges that are associated with INPDR are typically not known to the University of Melbourne IdP managers (or indeed to remote identity providers that may support authorisation information related to access and use of the INPDR). Whilst it is possible to embed these roles into the IdP, e.g. by populating attributes designed for this purpose (*eduPersonEntitlement*), the scale of projects occurring at the University of Melbourne would make this centralised source of attributes unworkable. Other more flexible solutions are therefore required.

The application of semantic web technologies in access control and trust management has drawn considerable attention in recent years [25]. In terms of access control model specifications, ontologies can play the role of formalising concepts and their inter-relationships. By combining OWL ontologies and RBAC models, Finin et al. proposed the *RBACOWL* and showed how a “reasoner” could establish access control decisions. Furthermore, hierarchies existing in roles/attributes may lead to policy inheritance, which can facilitate security management [26]. Based on this, Javanmardi et al. presented a semantic-based access control model (SBAC) in which all concepts in the subject, object and action domains could be formally defined [27]. By reasoning about subsumption in these domains, potential authorisations could be achieved by policy propagation. Since ontologies are a well-structured tool for knowledge construction, they can be used to share knowledge and avoid ambiguity. Leithead et al. exploited ontologies in credentials modelling [28]. By sharing such credentials between independent entities, they showed how problems such as information leakage and unnecessary disclosure in trust negotiation could be resolved effectively. In this paper, we apply ontologies to access control systems not only when there is a well-formed knowledge, but by reasoning and querying to achieve further potential authorisations based upon derived (reasoned) knowledge.

### III. AUTHORISATION AND TRUST MANAGEMENT IN THE INPDR

Through the INPDR, clinicians from around the world that deal with patients with Niemann-Pick Diseases must request to be enrolled to access and use the registry. Once ratified, they are assigned a username/password and a set of given privileges. Having enrolled, these clinicians are able (subject to privilege) to input a fixed set of pre-agreed patient data through targeted forms, which is subsequently used (again subject to access control privileges) for searching and analysing these data sets. Centres from around the world currently delegate the responsibility for assignment and management of privileges and their subsequent enforcement to the centralised INPDR provider. However even with this delegation, it is the case that all sites can be autonomous to varying degrees and define their own degrees of data sharing policies. Current data sharing

levels include the “*Centre*”, i.e. only researchers/clinicians from that Centre can access this data; the “*Country*”, i.e. only researchers/clinicians from that Country can access this data, and “*ALL*” where anyone who can authenticate to the INPDR registry is allowed to access this data. With this centralised model of delegation to the registry, explicit centralised trust on the assignment and enforcement of all aspects of security is adopted.

#### A. Role Definitions

Within the INPDR several pre-agreed roles have been identified as the basis for making access control decisions within the registry: “*Clinician*”, “*Local Collaborator*”, “*Local Researcher*” as well as “*Other Researcher*”. It should be noted that these roles were primarily identified through experiences in supporting a range of other similar registries and not through explicit requirements and needs that were identified by the INPDR research community. Based on the assignment and possession of these roles, people with different roles are subsequently able to perform different sets of actions on specific datasets. Specifically the roles and the allowed actions are shown in TABLE I.

TABLE I ROLE MEANING OF THE REGISTRY

Role Names	Local/Remote	Meaning
Clinician	Local	can create/edit/delete data for his/her Centre; can read all data tagged as his/her Centre; can read all data tagged as his/her Country; can read all data tagged as ALL;
Local Collaborator	Local	can create/edit data for his/her Centre; can read all data tagged as his/her Centre; can read all data tagged as his/her Country; can read all data tagged as ALL;
Local Researcher	Local	can read all data for his/her Centre; can read all data tagged as his/her Country; can read all data tagged as ALL;
Other Researcher	Remote	can read all data tagged as ALL;

It should be noted that all patients are associated with a specific Clinician whose email address is given (and the only truly identifying data in the registry). As well as being the only person where further information on the patient can be obtained or access to bio-specimens from the patient, the person with the Clinician role is responsible for ensuring that patient consent is obtained and ethics approval for entering data into the registry.

#### B. Access Levels of Patient Data

Once a patient record is created, a unique patient identifier is automatically generated by the registry, which consists of the NPD type, the standardised abbreviated country-code, the abbreviated clinic identity and an integer that is associated with that individual, e.g. “NPAB-AUME1-2 indicates the 2<sup>nd</sup> Niemann-Pick patient (with NPD type A/B) from Australia, Melbourne centre 1. Figure 1 shows a subset of patient listing in the INPDR.



NPA/B Patient Listing							Enrol new NPA/B Patient
Patient ID	Access level	Date of birth	Gender	ASMD type	Consent	Physician	
NPAB-AUMEI-2	All	08/02/2010	Male	A	Yes	email	Details
NPAB-ITUD2-7	Centre	11/08/1980	Female	B	Yes	email	Details
NPAB-ITUD2-7	Centre	11/08/1980	Female	B	Yes	email	Details
NPAB-UKBIS-4	Country	08/07/1993	Male	A	Yes	email	Details
NPAB-UKBIS-6	All	02/02/2010	Male	B	Yes	email	Details
NPAB-UKBIS-1	All	08/02/2010	Male	A	Yes	email	Details
NPAB-UKBIS-3	Centre	01/11/1996	Female	B	Yes	email	Details
NPAB-UKBIS-5	All	02/02/2010	Male	A	Yes	email	Details

Figure 1 Summary of NPD Type A/B Patient RECORDS

The INPDR is a completely centralised solution. It is currently used by clinicians/researchers from across the world including contributors from Australia, USA and across Europe. All sites need to be registered beforehand. The Clinicians need to be associated with those sites. The enrolment processes is based upon review of each individual case for access/enrolment. This has obvious scalability issues. When a clinician or researcher logs into the system, the security policy automatically restricts/enforces their abilities in the registry. Thus for example, creating a patient record in the system will have a patient record that is automatically populated with the Clinician and Centre related information based on the assigned privileges.

At present the system works relatively well given the number of collaborators and sites. However it will face scalability and other non-technical issues moving forward. Having a single source of authentication and authorisation information will be challenged where there are hundreds of sites and thousands of researchers involved. Furthermore, this centralised model cannot handle the decentralised evolution of organisations and people. Thus if a clinician moves hospitals or has a different role within the hospital then there is no way to identify this. An improved model would be to support decentralised security authentication and utilisation of existing/distributed authorisation credentials.

Thus instead of setting one central registry used to authenticate all users and define and assign roles/attributes for each user, we assume a decentralised model whereby authentication is made by remote sites, and authorisation information (credentials) are asserted by one or more remote sites and used (reasoned about) to make local enriched access control decisions. However this raises a range of challenges. Federated authentication is dependent upon standardisation and agreement of a range of pre-specified/agree core attributes and degrees of trust. Within the AAF this is based upon the SAML2 protocol [29] and exchange of *eduPerson* attributes with an underpinning PKI to ensure secure communications between the IdPs and SPs in the federation. Such federated access control does not exist across hospitals and certainly not (yet!) in an international setting.

A second challenge is in the attributes used for finer-grained access control decisions and their lack of standardisation. In an international setting, language-specific roles will be commonplace and will not work with existing security systems expecting the roles of “clinician” etc. A multitude of sources of authority will exist with keys and processes that they use for assignment of keys that are used for signing trusted attributes. Semantic reasoning has the ability to overcome many of these issues. The Shibboleth-based

architecture that we have adopted to illustrate the utilisation of semantic web technologies in this context is shown in Figure 2. Compared with the INPDR’s built-in centralised system taking the responsibilities of authentication and authorisation, this decentralised model leverages several trusted sites through delegating a set of pre-agreed roles, attributes and privileges.

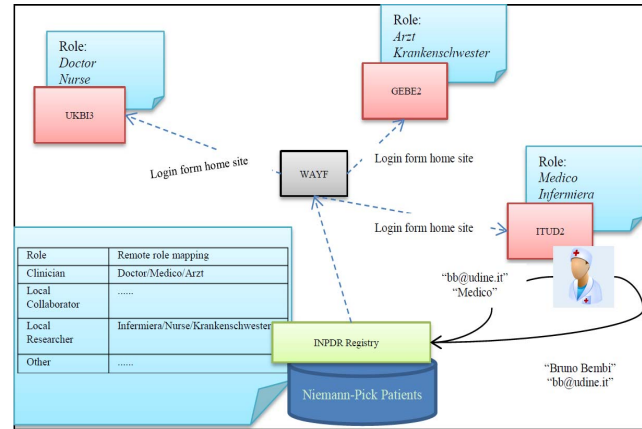


Figure 2 Underpinning decentralised access control model

For example, Dr. Bruno Bembi working for the “ITUD2” (Hospital-2 from Udine, Italy) may send a request to the INPDR for records of patients with NPD. With decentralised authentication, he will be redirected back to his home site for identity authentication. The Where-Are-You-From (WAYF) service is a core component in Shibboleth systems used for selection of IdPs. In this scenario, Dr. Bembi selects ITUD2 and confirms his identity. At this point several possibilities exist. The ITUD2 IdP may choose to send the authentication information and roles (attributes) that they are willing to release to the INPDR. More complex negotiations are also possible at this point depending on the level of trust. We assume here that his role as a doctor (*Medico*) is released as a digitally signed and trusted attribute certificate and that his identity is recognised by the INPDR as a trusted collaborator. We also assume that such international federations are possible and supported. In the existing INPDR system, such a vocabulary would not allow access. However semantic web provides far richer reasoning capabilities.

#### IV. ONTOLOGY-BASED SECURITY FRAMEWORK

By leveraging federated authentication and re-use of existing privilege related information, it is no longer necessary for the registry to be solely responsible for pre-assignment of static sets of “identity-role/attribute” mappings. As such, problems about scalability can be resolved to some degree. Thus if a doctor leaves a given hospital they will no longer be able to authenticate at that hospital and hence federated access to the registry will not be possible. However, this model also raises other problems in terms of semantic understanding. In the healthcare context, it is often the case that attributes with identical names issued by different organisations may have different meanings, while conversely other attributes with the same semantic meaning may be assigned different literal names. For example, to obtain a “doctor” certificate in Canada it is necessary to pass the Medical Council of Canada Evaluating

Examination (MCCEE. <http://www.mcc.ca/en/exams/>). This examination contains completely different programs compared with the Australian Medical Council examination (AMC. <http://www.amc.org.au/>). This means in some cases, “Doctor” credentials issued by different sites should be granted with different privileges to certain datasets. Furthermore language issues also arise, e.g. in the UK/US/Australian terms such as *doctor*, have international counterparts such as French (*docteur*), German (*arzt*) and Italian (*medico*). These relationships cannot be matched by standard security approaches. The same issues apply to nurses, French (*infirmier*), German (*krankenschwester*), and Italian (*infermiera*), as well as for other more specialised roles, e.g. *paediatric endocrinologist*. For these issues, semantic-based approaches such as ontology engineering can play the role of a “bridge” between vocabularies, which may be “different” literally and/or semantically.

#### A. Formalising the Access Control Model

To illustrate the potential of semantic web technologies in the e-Health security domain we explored a range of case studies around key challenges associated with the INPDR. We illustrate the realisation of these challenges using the ontology editor, protégé 4.0 (<http://protege.stanford.edu/>) to represent common elements by “class” (e.g. roles, subjects, countries, etc.) and specific instances by “individual”, e.g. “Clinician” under the class of Role and “Udine02” under the class of Italy. Such concepts can be further related by different properties such as “belongTo”, “hasRoleIs” and “accessLevelIs” etc. For example in Figure 3, as one instance of class “PatientData”, the record of “NPAB-UKBI3-1” is set with three properties: “belongTo” the centre Birminham03, the “accessLevel” as Centre and the “typeIs” NPAB (one form of Niemann-Pick disease).

Property assertions: NPABUKBI0301	
Object property assertions	+
belongTo	Birmingham03
Data property assertions	+
accessLevelIs	Centre
typeIs	NPAB

Figure 3 Specification of patient “NPAB-UKBI3-1”

With these basic concepts and relations, different enhanced types of authorisation can be described and enforced as a “restricted class”. Authorisation at different levels can be realised by “linking” the restricted subject and object in pairs. For example, the authorisation of viewing data for the Birmingham03 can be divided into two parts: retrieving the target data and filtering the qualified users as shown in Figure 4 and Figure 5. In this way, improved reusability of the discrete classes can be achieved. Rather than explicitly code the security policy, flexible reasoning about the policy using pre-existing data is possible. In this model, access to read data requires both a translation of the credentials that the user may

have and the potential data sets that may exist that can be satisfied by presentation of these credentials.

```
<owl2xml:ClassAssertion>
<owl2xml:Class owl2xml:URI="&OntologyCaseStudy;CanbeViewedByBI03"/>

<owl2xml:ObjectIntersectionOf>
<owl2xml:Class owl2xml:URI="&OntologyCaseStudy;PatientData"/>
<owl2xml:ObjectSomeValuesFrom> /*ObjectProperty*/
<owl2xml:ObjectProperty owl2xml:URI="&OntologyCaseStudy;belongTo"/>
<owl2xml:ObjectOneOf>
<owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;Birmingham03"/>
</owl2xml:ObjectOneOf>
</owl2xml:ObjectSomeValuesFrom>

<owl2xml:DataHasValue> /*DataProperty*/
<owl2xml:DataProperty owl2xml:URI="&OntologyCaseStudy;accessLevelIs"/>
<owl2xml:Constant>Centre</owl2xml:Constant>
</owl2xml:DataHasValue>
</owl2xml:ObjectIntersectionOf>
```

Figure 4 The data can only be viewed for Birmingham03 (centre only)

```
<owl2xml:Class owl2xml:URI="&OntologyCaseStudy;CanViewBI03"/>

<owl2xml:Class owl2xml:URI="&OntologyCaseStudy;Subjects"/>
<owl2xml:ObjectUnionOf>
<owl2xml:ObjectSomeValuesFrom> /*ObjectProperty*/
<owl2xml:ObjectProperty owl2xml:URI="&OntologyCaseStudy;hasRoleIs"/>
<owl2xml:ObjectOneOf>
<owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;Clinician"/>
<owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;LocalCollaborator"/>
<owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;LocalResearcher"/>
</owl2xml:ObjectOneOf>
</owl2xml:ObjectSomeValuesFrom>
</owl2xml:ObjectUnionOf>

<owl2xml:ObjectSomeValuesFrom> /*ObjectProperty*/
<owl2xml:ObjectProperty owl2xml:URI="&OntologyCaseStudy;belongTo"/>
<owl2xml:ObjectOneOf>
<owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;Birmingham03"/>
</owl2xml:ObjectOneOf>
</owl2xml:ObjectSomeValuesFrom>
```

Figure 5 Subjects can view the data only for centre “Birmingham03”

To further represent the enhanced function offered by semantic web for reasoning and querying, consider users (X, Y and Z) with associated information as shown in TABLE II. We consider a scenario where user X has interest in the information related to patient “NPABUKBI3-01”. To determine if he/she has the right to view the patient data, the relevant policy rules must be evaluated. Firstly, by querying with the patient ID, only one rule can be filtered out, since the patient from “Birmingham” has the access level set as “Centre”. At this stage, the corresponding querying of the “link” relation is checked to establish all qualified subjects. If the current requestor is included, then the privilege will be granted. As shown in Figure 6, since the result is matched, user X will be granted access to view the record.

TABLE II EXAMPLES OF REQUESTORS

	Country	Centre	Role
X	UK	Birmingham	Clinician
Y	Australia	Sydney	Other Researcher
Z	Italy	Udine	Local Collaborator

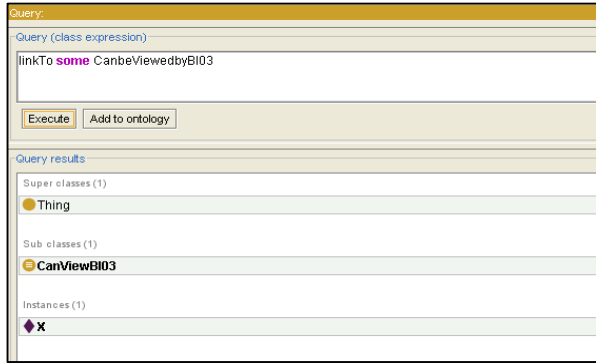


Figure 6 Checking the qualified requestors

## B. Discussion

### 1) Enhanced Access Control through Federation

Normally a federation consists of a group of organisations that are willing to share data and services (and hence trust one another). Through authenticating at their home sites, pre-agreed roles/attributes can be assigned to users for further role matching. A key challenge is the trust to initially enter the federation. Most identity federations are based upon established levels of trust, e.g. trust all universities on their processes for local identity management. In the e-Health domain, further fine-grained checks on the trust of an organisation will be needed. For example the reputation of the organisation and a much more rigorous check on the identity management procedures, e.g. do they have strong password protection policies and are users expected to change the passwords regularly. Assuming these checks are passed, the federation may send an invitation to the organisation for it to be a federated member [33]. Joining the federation provides a minimum of trust since it provides trusted authentication and does not yet include the necessary privileges to access sensitive resources such as the INPDR. Nevertheless, this model of federated authentication is essential to support flexible and scalable solutions based around systems such as INPDR since the centralised registry is completely ignorant of the local authentication systems and user management practices, and indeed whether the user is still a member of the given organisation.

### 2) Enhanced Role Interpretation

As mentioned, access to and use of INPDR requires deeper levels of trust. Newly added centres require further negotiations in order to gain deeper trust. In any given federation, there should be a pre-agreed set of attributes and relations that exist. These should be used as the basis for any negotiation with regard to deeper trust. Consider a scenario where a German Medical Certification Department can issue credentials to German clinicians (in German “Arzt”). Without any annotation, the INPDR cannot recognise the information because it is in a different literal form “Clinician”. However, through accepting the practices used for its certification program, stakeholders of the INPDR may accept this role as a Clinician when dealing with the requests. Therefore, they agree to the mapping between “Arzt” and “Clinician” (see the mapping description in Figure 7). As shown in Figure 8, by releasing the role of *Arzt*

from local organisation, user K can request to view the record of patient “NPAB-GEBE01-7” (See attributes “Country: Germany”, “Centre: Berlin01” and “Access Level: Centre”) of Figure 9.

```
<owl2xml:SameIndividuals>
<owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;Clinician"/>
<owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;Arzt"/>
</owl2xml:SameIndividuals>
```

Figure 7 Setting the role mappings

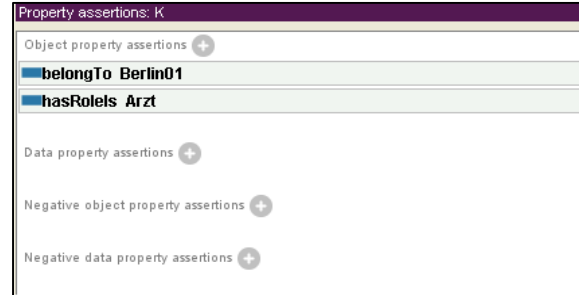


Figure 8 Assertion of attributes of requestor K

Since it can be reasoned that the user is a Clinician from Germany, the authorisation result shows user K can view the record in that centre (authorisation of viewing “NPABGEBE01-7” is evaluated by “CanViewBE01”).

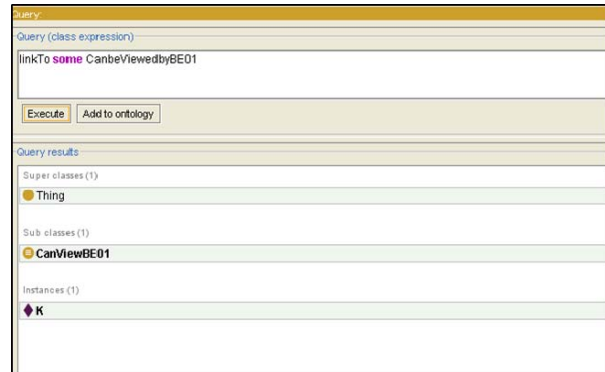


Figure 9 Checking the qualification of user K

To show the function of improved semantic understanding, equivalent mapping between terms can be undertaken. As the ontology supports inference, such relations can be established based on richer analysis about the concepts themselves. Through making use of existing attributes, semantic models can infer other properties of policies. For example, some hospitals may not have the “Centre” attribute given explicitly in the user’s personal description; however such information can be ascertained from the email address.

### 3) Enhanced “Country” and “Centre” Reasoning

Semantic web allows expressing basic concepts and relationships to make authorisation decisions with incomplete information. For example, consider patient “NPAB-UKBI03-4” restricted for users with the access level of “Country”. Any user from a centre in the UK (Birmingham, London ...) can access this patient detail even without explicitly providing their originating country. As shown in Figures 10-12, without any

specification of nationalities, it is possible to infer that a user from centre London01 can see “NPAB-UKBI03-4” as London is part of the UK.

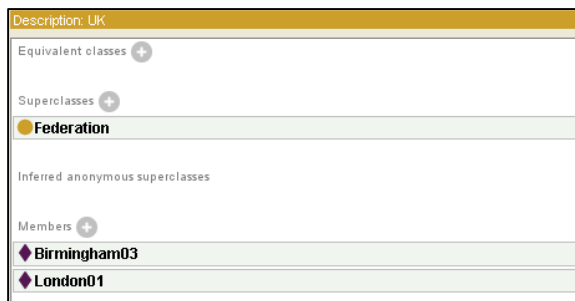


Figure 10 Federation structure of UK Organisation

```

<owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;X1"/>
<owl2xml:ObjectPropertyAssertion>
<owl2xml:ObjectProperty owl2xml:URI="&OntologyCaseStudy;belongsTo"/>
  <owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;X1"/>
  <owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;London01"/>
</owl2xml:ObjectPropertyAssertion>
<owl2xml:ObjectPropertyAssertion>
<owl2xml:ObjectProperty owl2xml:URI="&OntologyCaseStudy;hasRoleIs"/>
  <owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;X1"/>
  <owl2xml:Individual owl2xml:URI="&OntologyCaseStudy;Clinician"/>
</owl2xml:ObjectPropertyAssertion>

```

Figure 11 Attribute assertions for requester X1

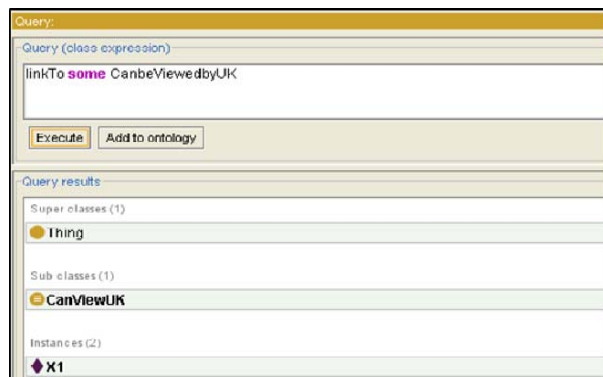


Figure 12 Enhanced reasoning between centres and countries (Attributes)

#### 4) Enhanced Role Hierarchy Scalability

Since the credential “Arzt” can be mapped to “Clinician” it should also relate to the role hierarchy of a Clinician. This can help in role hierarchy scalability. In this scenario, “Clinician” can perform all actions (delete/create/edit/read) on the specific data, while the “Local researcher” can only read these data resources. Someone with the German role of “Arzt” delegated as “Clinician” implies that he/she will be able to enjoy the privileges as given in that hierarchy. Without creating new contents in the original structure, it is possible to extend this by semantic bridges.

### V. CONCLUSIONS AND FUTURE WORK

Previous work on tackling authorisation and trust management in distributed systems has shown the limitations in terms of system scalability and autonomy. Inspired by the

opportunities offered by federated authentication, we propose a decentralised access control model with distributed agents helping in authenticating and issuing credentials. In this context, semantic web technologies can facilitate the policy specification and evaluation. We have formalised an ontology-based access control framework and explored its utility in the context of an international disease registry (INPDR). Based on a range of scenarios, we have shown how enhanced reasoning and querying can be achieved utilising semantic technologies. This work is based on the premise of a degree of trust underpinning the federated interactions, e.g. a PKI exists that is used for signing and hence trusting the security assertions. Such trust and hence federation does not yet exist however the technologies for federated authentication are now well established. Establishment of international trust federations is currently an active topic with work on integrating a range of national/international federated authentication systems.

The next phase of this work is to further explore strategies for negotiating trust and when to exchange sensitive information and with whom. More explicit models of roles and a suitable ontology applicable to e-Health environments will also be undertaken. The notion of reputation will be explored as a model for an on-going extension and refinement of existing trust mechanisms.

### ACKNOWLEDGMENT

The authors would like to thank the INPDR research community. The INPDR is supported by a grant from the European Union Directorate General for Health and Consumers (DG-SANCO) and a range of clinical partner organizations.

### REFERENCES

- [1] R. L. Richesson and J. E. Andrews, “Introduction to Clinical Research Informatics”, Clinical Research Informatics, Springer, London, 2012, pp. 3-16.
- [2] R. L. Richesson and K. Vehik, “Patient Registry”, Clinical Research Informatics, Springer, London, 2012, pp. 233-252.
- [3] Shibboleth2 Internet. <https://shibboleth.net/>. Accessed on 20<sup>th</sup> March.
- [4] Niemann-Pick Disease Overview – Type A, B and C, [http://www.nnpdf.org/npdisease\\_01.html/](http://www.nnpdf.org/npdisease_01.html/). Accessed on 22<sup>th</sup> March.
- [5] E. H. Schuchman, “The pathogenesis and treatment of acid sphingomyelinase-deficient Niemann-Pick disease”, International journal of clinical pharmacology and therapeutics, vol. 47, 2008, pp. 48-57.
- [6] J. Watt, and R.O. Sinnott, “Supporting Federated Multi-Authority Security Models”, Proceedings of the 11<sup>th</sup> IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, Newport Beach, CA, USA, May 2011.
- [7] J. Wei, J. Arshad, and R.O. Sinnott, “A Review of Grid Authentication and Authorisation Technologies and Support for Federated Access Control”, Journal ACM Computing Surveys, vol. 43, issue 2, January 2011.
- [8] T. Berners-Lee, J. Hendler and O. Lassila, “The semantic web”, Scientific American, vol. 284, issue 5, 2001, pp. 28-37.
- [9] A. Maedche, “Ontology learning for the semantic web”, Springer Science & Business Media, 2002.
- [10] D. Hardt, OAuth 2.0 Authorisation Framework, <https://tools.ietf.org/html/rfc6749/>. Accessed on 20<sup>th</sup> March.
- [11] S. T. Kent, “Internet privacy enhanced mail”, Communications of the ACM, vol. 36, issue 8, 1993, pp. 48-60.
- [12] S. Garfinkel, “PGP: pretty good privacy”, O’Reilly, 1995.
- [13] R. Sulaiman, D. Sharma, W. Ma, and D. Tran, “A security architecture for e-health services”, 10<sup>th</sup> International Conference on Advanced Communication Technology (ICACT), vol. 2, 2008, pp. 999-1004.



- [14] G. Russello, C. Dong, and N. Dulay, "A workflow-based access control framework for e-health applications", 22<sup>nd</sup> International Conference on Advanced Information Networking and Applications-Workshops, 2008, pp. 111-120.
- [15] P. Ruotsalainen, "A cross-platform model for secure Electronic Health Record communication", International Journal of Medical Informatics, vol. 73, issue 3, 2004, pp. 291-295.
- [16] M. Humphrey, J. Basney and J. Jokl, "The case for using Bridge Certificate Authorities for Grid computing", Software: Practice and Experience, vol. 35, issue 9, 2005, pp. 817-826.
- [17] M. D. Abrams and A. B. Jeng, "Network security: Protocol reference model and the Trusted Computer System Evaluation Criteria", IEEE Network Mag., vol. 1, issue 2, 1987, pp. 24-33.
- [18] S. Osborn, R. Sandhu, and Q. Munawar, "Configuring role-based access control to enforce mandatory and discretionary access control policies", ACM Transactions on Information and System Security (TISSEC), vol. 3, issue 2, 2000, pp. 85-106.
- [19] H. Lindqvist, "Mandatory access control", Master's thesis, Umea University, Department of Computing Science, 2006.
- [20] E. Bertino, P. A. Bonatti, and E. Ferrari, "TRBAC: A temporal role-based access control model", ACM Transactions on Information and System Security (TISSEC), vol. 4, issue 3, 2001, pp. 191-233.
- [21] S. M. Chandran and J. B. Joshi, "LoT-RBAC: a location and time-based RBAC model", In Web Information Systems Engineering-WISE 2005, Springer, Berlin, Heidelberg, 2005, pp. 361-375.
- [22] F. Hansen & V. Oleshchuk, "SRBAC: A spatial role-based access control model for mobile systems", In Proceedings of the 7th Nordic Workshop on Secure IT Systems (NORDSEC'03), 2003, pp. 129-141.
- [23] O. O. Ajayi "Dynamic trust negotiation for decentralised e-health collaborations", Doctoral dissertation, University of Glasgow, 2009.
- [24] Australian Access Federation (AAF). <http://www.aaf.edu.au/>. Accessed on 24<sup>th</sup> March.
- [25] R.O. Sinnott, T. Doherty, N. Gray and J. Lusted, "Semantic Security: Specification and Enforcement of Semantic Policies for Security-driven Collaborations", Proceedings of 7<sup>th</sup> HealthGrid Conference, Berlin, Germany, July 2009.
- [26] T. Finin, A. Joshi, L. Kagal, J. Niu, R. Sandhu, W. Winsborough, and B. Thuraisingham, "ROWLBAC: representing role based access control in OWL", In Proceedings of the 13th ACM symposium on Access control models and technologies, 2008, pp. 73-82.
- [27] S. Javanmardi, M. Amini, and R. Jalili, "An access control model for protecting semantic web resources", In Web policy workshop, 2006.
- [28] T. Leithhead, W. Nejdl, D. Olmedilla, K. E. Seamons, M. Winslett, T. Yu and C. C. Zhang, "How to Exploit Ontologies for Trust Negotiation", In ISWC Workshop on Trust, Security, and Reputation on the Semantic Web, vol. 127, 2004.
- [29] P. Mishra, "SAML V2.0 X.500/LDAP Attribute Profile", 2006.
- [30] Medical Council of Canada, <http://www.mcc.ca/en/exams/>. Accessed on 20<sup>th</sup> March.
- [31] Assessment & Examination. Australian Medical Council Limited, <http://www.amc.org.au/>. Accessed on 20<sup>th</sup> March.
- [32] Protégé 4.0, <http://protege.stanford.edu/>. Accessed on 20<sup>th</sup> March.
- [33] P. A. Cabarcos, F. Almenárez, F. G Mármol and A. Marín, "To federate or not to federate: a reputation-based mechanism to dynamize cooperation in identity management", Wireless Personal Communications, vol. 75, issue 3, 2014, pp. 1769-1786.