An Infrastructure for Secure Sharing of Medical Images between PACS and EHR Systems

Kamran Sartipi, Krupa A. Kuriakose, and Weina Ma
Department of Electrical, Computer and Software Engineering
University of Ontario Institute of Technology
Oshawa, ON, L1H 7K4, Canada
{*Kamran.Sartipi, Krupa.Kuriakose, Weina.Ma*}*@uoit.ca*

Abstract

New advances in information and communication technologies (ICT) and their incorporation into the medical domain have created opportunities to enhance medical services and provide improvement to workflow at a low cost. However, to implement such services, the current medical system needs to be integrated, secured, and available to health professionals and patients. In this paper, we propose an infrastructure that suggests the use of techniques and standards such as: cooperative multi-agents, standards for user authentication and service authorization, as well as protocols for cross-enterprise document sharing. The proposed infrastructure allows for integration of a PACS (Picture Archiving and Communication system) with a widely accepted HL7 (Health Level Seven) standard infrastructure for provisioning nation-wide electronic health records (EHR). In this approach, the cooperative agents provide: i) an action-based access control mechanism to share medical images that allow safe integration of a PACS and the Diagnostic Image Repository (DI-r) systems within a standard EHR system; and ii) a behavior-pattern based security polity enhancement to assist the system administrator. Such secure and interoperable medical imaging systems are easy to expand and maintain.

Acknowledgement. This research was conducted with collaboration of Dr. David Koff and Dr. Peter Bak and Ms. Jane Castelli at MIIRC@M Centre of McMaster University and Mr. Duane Bender and Mr. Arun Agrawal at Mohawk College. This research was funded by an ORF grant for the project "Secure Intelligent Content Delivery System for Timely Delivery of Large Data Sets in a Regional/National Electronic Health Record".

1 Introduction

PACS are legacy systems that are used for storing and retrieving medical images. To restrict the privacy breaches or prevent intrusion from outside, the functionality of the existing PACS are localized to their working environment. The existing PACS are from different vendors with different designs for data scheme and communication workflow, which make them vulnerable against security attack and privacy violation. The limitations of the current PACS are: i) user authentication is restricted to the local system by using local user ID and password; ii) access control rules are local to their system; iii) patient identity is presented in different ways; iv) patient consent directives are not regularly used; and v) audit trail facility is localized to each system. Hence, a safe integration and communication of PACS and DI-r systems is a challenging task.

As a solution to the security issues in the exiting PACS, we propose a common infrastructure for secure sharing of medical images between multiple PACS and DI-r systems. Health Information Access Layer (HIAL) [3], which lies between multiple PACS and the DI-r systems acts as a mediating layer for communication and data transfer. Since HIAL supports various interoperability standards and technologies, it is an ideal platform for communication.

The proposed solution satisfies the system access control requirements by authenticating the PACS users with the help of a third-party identity provider [7] to access data from multiple image providers. Authorization [9] is granted by comparing the user's requested operation with the system's security policies, by considering patient consent directives, and by the audit trail [11] records. Furthermore, we extract the user behaviour [27] as patterns in sequences of user operations in an extended period of time, which are used for updating system security policies.

The major contributions of this paper are as follows: i) proposing an interoperable and secure architecture, which integrates with the conceptual architecture for EHR systems and is endorsed by HL7 and DICOM (Digital Imaging and Communications in Medicine) standards; ii) incorporating the most advanced privacy and security protocols to provide a federated identity management system; and iii) introducing a behavior-based technique which allows to enhance the system's access control policies.

The remainder of this paper is organized as follows: Section 2 describes the work that is relate to our approach. In Section 3 we discuss the underlying technologies used in the architecture. Section 4 presents the proposed architecture for secure image sharing. Section 5 includes two case studies. Finally, Section 6 provides a brief discussion on the proposed approach and conclusion.

2 Related work

Security and access control

There are different access control models to manage user authorisation. *Role Based Access Control* [26] is the most common method used. Each user is associated with a particular role and based on the assigned role the user can access the system. *Team Based Access Control* [13] allows a member of a team to enjoy the privileges and rights common for the team. *Content Based Access Control* [25] puts restrictions on the user access based on the content of the resource. *Attribute Based Access Control* [26] is based on the attribute or characteristics of the user. For the user to meet the access control requirements he needs to prove that his attribute can meet the policy of the system. In *Situation Aware Access Control* [28], assigning a role to a user and granting permission to that role is determined based on the situation of the system at that instant. Situation is defined by considering the previous device interactions and the variation of a relevant set of context related to a specific application. *Scenario Based Access Control* [26] proposes various steps to define tasks, and the user access right is determined for each step. *Context Aware Access Control* [26] examines the context of the user at a particular instant and the user is granted access permission based on the context.

Federated Identity Management

Federated Identity Management (FIM) is an identity management solution that allows users to authenticate themselves with an infrastructure that has common sets of access policies and all participants share a common trust. Once authenticated with this structure, the user can securely access data from a third party that is registered to FIM infrastructure [16]. In [22] Deng et al. propose a scheme to preserve patient identity in federated eHealth systems. They use a cryptographic algorithm that allows issuing context specific local identifiers to users. This local identifier is derived from a unique global identifier. Campos et al. [17] propose a centralized infrastructure for authentication and secure identity management for eHealth by making use of their eID smart-card. This can be done by establishing trust between Government and healthcare systems; the Government issues legal policies and the patient defines his/her consents. Peyton et al. [18] look into the business and technical issues in Liberty Alliance federated identity solution by using a simple ePrescription scenario. Canada Health Infoway (CHI) defines requirements and specifications for Security and Privacy Architecture for Canadian healthcare infrastructure [1]. CHI are conducting various research to develop interoperable healthcare systems that are compatible with the existing medical standards and associated communication technologies for medical data exchange. The CHI has also defined EHR Security and Privacy Requirements for healthcare domain and has highlighted the restrictions on medical data usage [4]. Even though a number of documents are addressing the requirements and policies, a real infrastructure has not yet been implemented for EHR. In our approach, we define a common infrastructure that allows PACS, with different functionalities and features, to share their

resources. This is possible by authenticating users with a single identity provider and then the access control module uses the credentials issued by the identity provider. This ensures secure data sharing among the systems.

Agents and healthcare

A software agent performs a task on behalf of a user or another program. It is mainly applicable in situations where a user is supposed to do a task as part of the system functionality, but unable to do so in an online manner. In such situations, the user employs an agent to do the task on his behalf. Different types of agents include: autonomous agents, intelligent agents, distributed agents, and mobile agents. Tian and Tianfield [24] studied the characteristics of healthcare delivery systems and encapsulated these functionalities into multiple agents, some of which include: *Personnel agent, Resource agent, Function agent*, and *System agent*. These agents play different roles and are integrated to perform the operations of a complete healthcare delivery system. Gupta and Pujari [14] propose a multi-agent solution for healthcare and medical diagnosis purposes, by collecting user input, translating them to knowledge and passing them to specialised agents for diagnosis and storing of the reports. Silva-Ferreira et al. [12] employ multiple agents to discover and retrieve patient information from multiple healthcare institutions using openEHR queries and HL7 messages to enable agents to query local repositories, retrieve patient details, and store them in the openEHR repository. Zhou [19] proposes a health evaluation approach based on multi-agents, including: simple reflex agents, goal-based agents, utility-based agents, and learning agents that are integrated using web services. Sulaiman, Huang and Sharma [23] propose a security mechanism for data communication by employing mobile agents. They use a multilayer communication approach and once the layer is chosen, appropriate mobile agents do the task of data transfer. In our work, we employ multiple agents to deal with access control decision making and an alerting system for the administrator to track user behavior. This includes an action agent to compare user's action with the system security policies; a consent agent to look into the patient consent directives and make appropriate decisions; and a behavior agent to notify the system's administrator regarding updating of policies.

3 Technologies

PACS Architecture

The architecture of a PACS is shown in Figure 1. The main components of a PACS are: image modalities, acquisition gateways, PACS controller and associated database and server, long term and short term archives, and workstations. A PACS is capable of acquiring, storing, transferring and retrieving medical images in healthcare environments [21]. PACS mainly rely on DICOM [20] and HL7 [15] standards for communicating with different image modalities (defined below). *Image modalities* are image acquisition components that capture medical images of patients. These include: X-ray, MRI, CT, fluoroscopy, etc. Some modalities capture images in digital format while others in analog format. For example, some X-ray scanners provide images in analog format. To deal with such situations PACS have an *acquisition gateway*, which is usually a computer system that is located between the image modalities and PACS environment. They convert analog images to digital format, thereby making it compatible with PACS. Acquisition gateway, if connected to Hospital Information System (HIS), can add additional information to patient images by using HIS interface and the HL7 protocol. *PACS controller* is the most important component of the system. It has multiple functionality as image storage is concerned. Images obtained from modalities are stored in the PACS database. When an image arrives, the text associated with the image is extracted; the image is compressed; the workstation to which the image has to be forwarded is determined; and the image is stored in an archive if it is not meant for immediate use. The PACS database, server and archiving systems are associated with the PACS controller. The PACS database is responsible for grouping and ordering of the images. It is connected to the *Radiology Information System* (RIS) to retrieve the data associated with the patient images. After properly arranging the images the recent ones are stored in the *short-term archive* and further to the *long-term archive* for future use. The user at his workstation views all the medical images. *Workstations* include software that support procedures for accessing images from the image database, processing images, and all user activities while working with medical images.

DI-r

The DI-r project initially started with eight

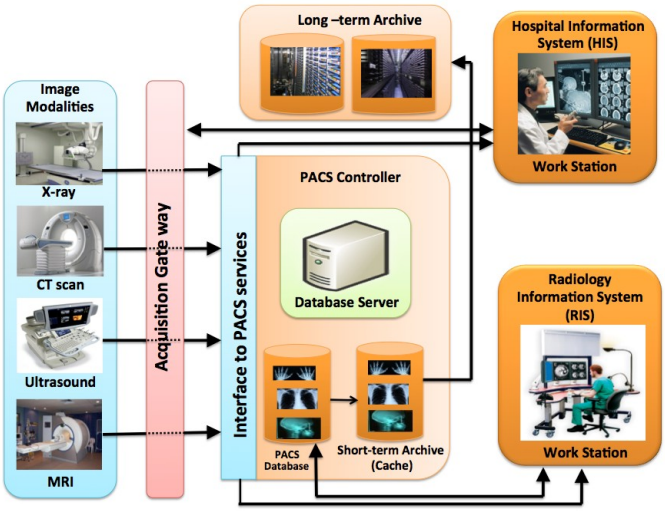

Figure 1: PACS architecture

hospitals replacing their film X-rays by PACS systems. Sanction is given to authorised medical personnel to share medical images securely with other members of a particular DI-r. The data stored is retrieved in the digital format, which makes it easier for communication. The system benefits the patients as well as the clinicians. From the patient's point of view, the DI-r reduces unnecessary travel, wait times in hospitals, repeated examination and after all reduces the number of times the patient has to be exposed to radiation. From the clinician's point of view medical images could be retrieved anytime from anywhere in the world. Further, it helps in faster diagnosis without wasting time for image recovery. In Ontario, the hospitals are partitioned into four clusters, each with a separate DI-r. These four DI-r clusters are: the Southwest Ontario Diagnostic Imaging Network (SWODIN) [10], the Hospital Diagnostic Imaging Repository Services (HDIRS) [6], the Northern and Eastern Ontario Diagnostic Imaging Network (NEODIN) [2], and the Greater Toronto Area West Diagnostic Imaging Repository (GTA West DI-r) [5]. Such DI-r clusters can further be integrated into a nation-wide document sharing infrastructure, which can also be integrated with a nation-wide EHR to provide full accessibility to medical images. There are a number of challenges for a fully functional infrastructure. The vendors are not yet compliant with an implementation of the imaging interoperability standards, namely "Cross-enterprise Document Sharing for Imaging" (XDS-I) and the "Integrating the Healthcare Enterprise Patient Identifier Cross Referencing" (IHE PIX). The Enterprise Master Patient Index (EMPI) should also be incorporated to achieve wide-scale interoperability. This ensures that each patient is represented only once across the imaging systems.

XDS-i

The Integrating the Healthcare Enterprise (IHE) organization released Cross-Enterprise Document Sharing for imaging (XDS-i) for medical image sharing. The shared document includes the results of imaging studies obtained from different image modalities and reports of image interpretation for the purpose of diagnosis [8]. The XDS-i protocol allows the user to coordinate activities related to locating and accessing images from a storage location. Medical images of a patient with different image modalities can be produced from any hospital or medical organization. For efficient retrieval of the images, a common location for storage of all medical images are used, where the image details and patient IDs are stored in different repositories. Figure 2 illustrates the steps of retrieving a medical image. The XDS document repository and XDS registry are intended to store necessary image information and register associated patient IDs. In a federated environment an authorized user queries the image details from the XDS document registry, which allows the user to

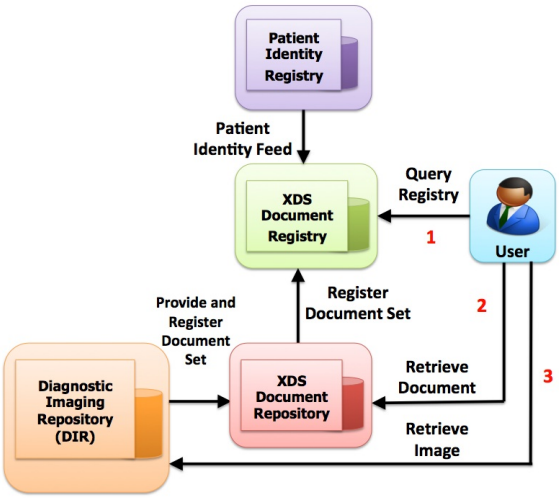

Figure 2: Cross-enterprise Document Sharing for imaging (XDS-i).

gather the image manifest from the XDS document repository; and consequently the image itself will be retrieved from the DI-r.

DICOM and HL7

DICOM is a standard developed by American College of Radiologists (ACR) and National Electrical Manufacturers Association (NEMA) for handling medical image communication between various image modalities and associated systems [20]. Since it includes the file format definition and network communication protocols, we can exchange medical images between two parties that support DICOM. It has the capability to store, query, and retrieve images from systems such as PACS. Its storage commitment feature confirms that an image has been permanently stored in a disk. Modality work list gives a description of patient details and inspection of patient examination schedule. DICOM also enables image modalities to keep track of the images and their details taken in that device. This helps the radiology department to keep track of the sequences of image transactions that is taking place on the device.

HL7 [15] is an organisation whose mandate is research for developing standards for healthcare interoperability of medical data. HL7 standards serve as the application layer (seventh layer) in the ISO – OSI (International Organization for Standardization- Open Systems Interconnection) model for communication. HL7 helps in standardising medical data exchange, it's management, and integration with other medical service providers. Different hospitals and healthcare providers use different format for data storage. In order to exchange this medical data among heterogeneous systems, HL7 provides a number of standards and methodologies which include conceptual standards, document standards, application standards and messaging standards.

OpenID and OAuth

OpenID and OAuth are decentralised open web standard protocols for authentication and authorization purposes in security, identity management and access control domains. Even though they can be used together in a system, their functionalities are different. OpenID reduces the need of multiple identities a user has to maintain while using different web-enabled applications and services. Whereas, OAuth provides a way to grant permission to a third party application to access the user's data on a server, without providing user's credentials the third party.

OpenID: With advances in web technology and increasing number of web applications that require registration of their users, the users are expected to take care of too many user-names and passwords [7], or use the same credentials for several services. On the other hand, the service providers are liable to keep their users' credential in secure databases, which may not always afford to do so. In either case, a single intrusion by a hacker can seriously affect the security of the user data. To solve the above problems, OpenID provides a single sign-on protocol to relieve both the users and the service providers. The main components of the OpenID authentication mechanism are as follows: *OpenID Provider* (IdP), a website that provides user specific URLs for authentication purposes; *Relying party* (RP), a system that requires to verify the user's authentication using the user's URL provided by IdP; and *User*, a person who desires to make use of services offered by the RP using the URL provided by the IdP. Once the user registers with the identity provider (IdP), he/she can login to all OpenID enabled web sites. IdP login allocates a URL to the user which contains a set of HTML tags used to identify the user's IdP.

OAuth: In legacy authentication models, the external applications could get access to the owners'

protected resources, meaning that there was no restriction on the duration or the amount of resources that they can access. Moreover, the resource owners had no provision to block a particular third party from access to their resources. The user had to deny all third party accesses by changing the password. OAuth addresses this issue by introducing an authorisation layer that separates the client from the resource owner. Whenever the client requests access to resource in the server, a token is issued to the client based on the authorisation grant from the resource owner. There are four main components in OAuth, as follows [9]: *Resource owner* grants access permission to the protected resources; *Resource server* hosts the protected resources; *Client application* is granted permission to access the contents in the resource server with the authorisation of the resource owner; and *Authorization Server* issues access tokens to the client application after successfully authenticating the resource owner and obtaining authorization. OAuth allows limited access (by a third party application) to a particular service, either through an access grant interaction between the resource owner and the targeted service, or by allowing a third party application to use the application on its own identity. The authorisation server, with the approval of the resource owner, issues access tokens to the third party clients. Later the clients use the issued tokens to access resources hosted by the server.

4 Proposed framework

In Figure 3, the large box in the middle illustrates the "EHR Infastructure" proposed by Canada Health Infoway [1], which constitutes the underlying EHR infrastructure in our approach, and the blue-background boxes around it represent different components for secure sharing of medical images between the PACS and DI-r systems. The main task of the EHR is to integrate all health related information into one infrastructure. This will help in secure medical data exchange between Point of Services (PoS) and various repositories that store the information. Authorised clinicians and healthcare providers can access these data based on the access right and system security policies. The communications between PoS and various repositories in EHR take place through HIAL (Health Integration Access Layer). HIAL acts as a gateway to separate the PoS from the data repositories. HIAL consists of several components, roles, and messaging standards to ensure interoperability when different systems are involved in data exchange. It has two layers of services: "Communication bus" services with communication capabilities responsible for exchange of messages and "Common bus" services that provide functionalities that are common to applications using the system.

We have enhanced the CHI Infastructure by adding components that resolve the security issues in sharing images between the PACS and DI-r systems. Additional repositories and registries support the major components in the architecture. The main objective is to take care of the authentication and authorization (access control) aspects of the system. *Access Control* component collects relevant user data and redirects the user for authentication. It then looks into granting access permission for the user. *Behaviour Agent* extracts the user behavior patterns and contacts the security administrator to deal with updating system policies based on user behavior.

Figure 4 illustrates an overall view of the proposed framework and emphasizes on the interactions between the components that are responsible for the secure sharing of images. When the user approaches the PACS system to retrieve an image from the DI-r, the "Access Control" component captures the relevant information of the user required for making access control decisions. We call this interaction an "Action" of the user. The attributes of these interactions are recorded according to an "Action Tuple" as follows:

Action Tuple = <User, Role, User Location, Server Location, Time of Day, Team, Delegation, Requested Profile Status, Service Invocation Type, Requested Data Type, Login/Logout Event, Emergency>

Once this information is gathered, the system assesses the user's credentials and the type of requested operation, to authorize or deny the requested operation. We also collect the attribute values of the sequence of user actions to extract the pattern of user's behavior which is used by the system administrator to adjust the system's policies for access control. In the following subsections, the Access Control and Behavior aspects of the proposed architecture are explained.

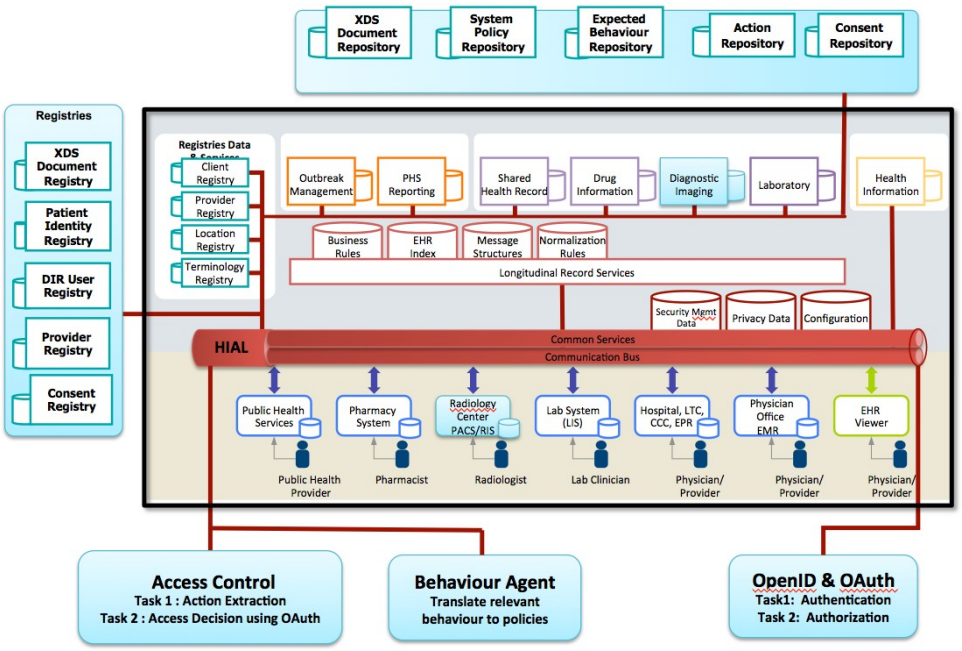

Figure 3: Overall view of EHR Architecture augmented by the proposed infrastructure.

4.1 Access control

The proposed architecture in Figure 4 shows that all access control decisions are centralized by the Access Control Component (ACC) which makes decision on behalf of the DI-r and XDS repositories (i.e., real Resource Providers). Figure 5 illustrates the details of the Access Control component and its associated components that together provide an advanced decision making mechanism to authorize access to images in the DI-r repository using the OpenID identify provider and the OAuth authorization procedure as defined in Section 3. The core of the Access Control component is the "Authorization Server" (AS). To make proper access control decisions, the AS must receive the information about: i) user's authentication; ii) nature of the access operation that complies with the system policies; iii) patient's consent directives; and iv) audit trail. The AS performs two main functions: authorization grant, and delivering the Access Token to the user service. Authorisation grant is given to the user service after authenticating the user and consulting with the Patient Agent (PA) and Action Agent (AA). The access control process is described below with reference to Steps 1 to 8 shown in Figure 4.

Step 1. The PACS user registers his/her credentials with an OpenID Provider that is trusted by both the user and the resource provider (i.e., DI-r component).

Step 2. The User-services unit (PACS interface service) registers the PACS as an image provider/viewer, in the DI-r Provider Registry.

Step 3. The User-services unit issues an "Access Request" to Access Control Component to transfer an image of a patient (e.g., retrieve from Cache/DI-r or store to Proxy to be scheduled for storing in DI-r).

Step 4. The ACC performs the authorization process using OAuth and OpenID protocols. It extracts the required information for authorization of the user-service, mentioned in 4(a) and 4(b) below. Then, by comparison of the extracted information ACC either grants or denies the requested access operation by the user. The details of the Access Control protocol are discussed in Subsection 4.1.

- *4(a) User Authentication*. The ACC obtains the web address of an OpenID provider that

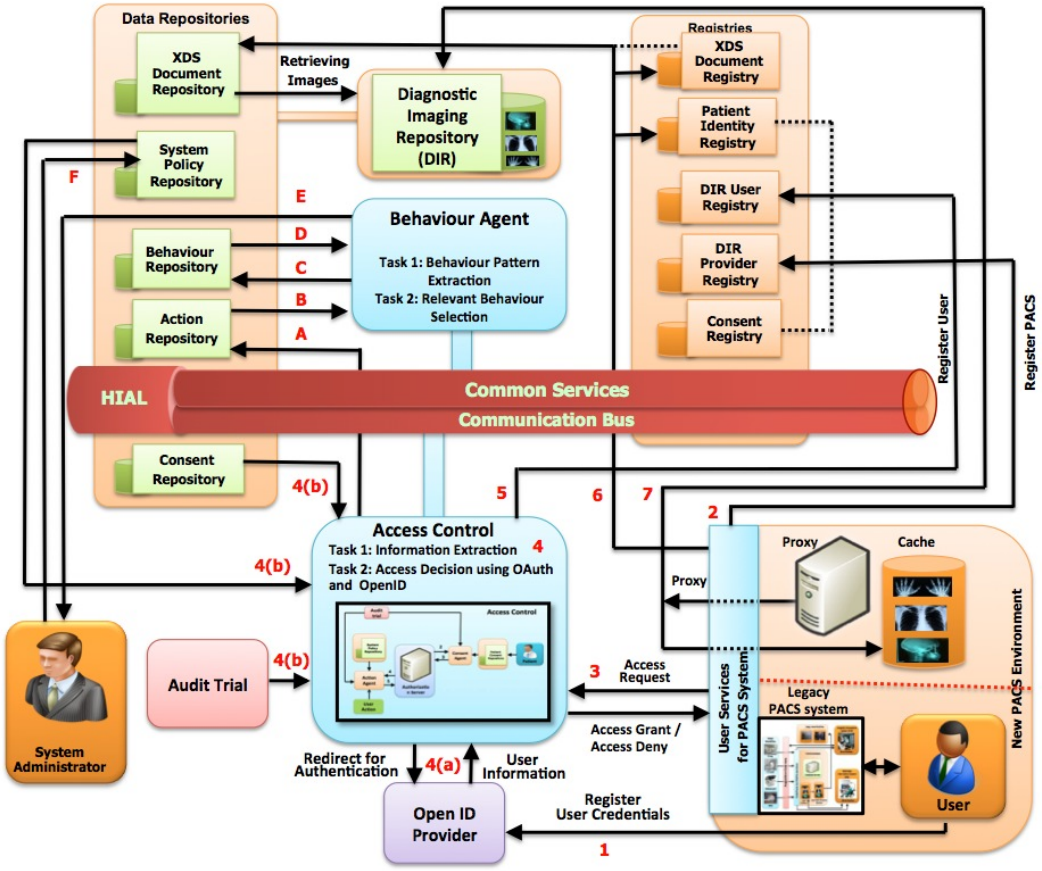

Figure 4: Proposed secure framework for PACS document sharing.

is trusted by both the user-service and image-provider service (DI-r/Cache/Proxy services). It then redirects the user-service to the OpenID provider web site for Authentication (userId and password). After authentication, the OpenID provider returns the required "user information" (not the credentials) to the ACC.

- *4(b) Information acquisition*. The ACC obtains the following information: i) type of access operation from access request message; ii) patient consent directives from "Consent Repository" component; iii) logged information from "Audit Trail" component; and iv) Access control policy rules from "System Policy Repository" component.

Step 5. After proper authentication of the user, in 4(a), the ACC registers the user in the "DI-r user Registry" as known user of the distributed PACS.

Step 6. After authorization of the requested access (Access Grant in Step 4), the user-service prepares for image retrieval or storage, as follows:

- *Retrieval*: it consults with the "Patient Identity Registry" and "XDS Document Registry" to determine whether the image of the intended patient is registered or not, and consults with the "XDS Document Repository" to determine the location of the image in the DI-r or Cache.

- *Storage*: the user-service registers the patient and the image in the above registries. Then it requests for a service to store the manifest of the image in the "XDS Document Repository".

Step 7. The user-service invokes a DI-r service to transmit the desired patient image from Cache

(immediately) or DI-r (with delay) to the PACS local storage (i.e., retrieval), or to move the image from the PACS local storage to the "Proxy" storage to be scheduled for transmitting to the DI-r at a proper time (i.e., delayed storage).

In the following part, the details of the ACC are described with reference to Steps A1 to A8 shown in Figure 5.

Step A1, *Access Request*. Access control operation begins with an access request message from the PACS user to the "Authorization Server" to transmit (retrieve or store) medical images between local storage of the PACS and the DI-r component. As the legacy PACS are proprietary, they lack proper APIs to integrate with other PACS and DI-r systems. This causes a major challenge to integrate these systems with advanced and standards based (HL7 and DICOM) systems. However, such an integration is inevitable for the future systems where the medical images will be shared by the PACS through nation-wide EHR systems. We assume that the legacy PACS will be equipped with proper and standard based "user services" (using reverse engineering techniques). Furthermore, the detailed information must be extracted from the access request message, including: *UserId, Role, Location, Time, Type of Operation, Requested Image, Emergency*, etc. These information will populate the Action Tuple discussed earlier in this section. The "User Action Extractor" in Figure 5 is an important module that captures these information from the network traffic.

Step A2, *Authentication*. Authorization Server (AS) presents a list of trusted Identity Providers (IdP) to the user and the user selects an IdP link that he has already registered with. After agreeing on a proper IdP, the AS and IdP establish a shared confidential code for that session and the AS redirects the user to IdP web site with an authentication request, where the user is authenticated using his/her OpenID credentials. Next, IdP redirects the user back to the AS with an identity assertion which includes an association handle. The AS validates the assertion using the association handle and shared confidential code.

Steps A3 & A4, *Patient Agent*. Once the user is authenticated, the AS should authorize the user service for the requested access. The AS (using the validated user ID) sends a request to the Patient Agent to verify the user's privileges for accessing the image of a particular patient (A3). The Patient Agent consults with the "Patient Consent Repository" and "Audit Trail" to assess the user. After this screening process the Patient Agent sends back the response on behalf of the patient (*Consented* or *Refused*) to the Authorization Server (A4).

Steps A5 & A6, *Action Agent*. In the case of a *Consented* response from the Patient Agent the AS contacts the Action Agent to investigate whether the user's requested action is authorized against the system's security policies (A5). The Action Agent compares the attributes of the extracted Action Tuple by the "User Action Extractor" with the system policies in the "System Policy Repository" and checks with the "Audit Trail" history to see if any unusual past situation exists or not. If the requested action is authorized, the action agent returns a positive response to AS (A6).

Steps A7 & A8, *Image Retrieve*. If both the Patient Agent and Action Agent approve the user's requested operation to the AS, the AS issues an "Access Token" to the user service. The user using the Access Token requests the protected patient's medical image from the DI-r repository using the XDS-i protocol described in Section 3 (A7). Finally, the DI-r retrieves the image and sends it to the user's local PACS to be viewed (A8).

To illustrate the access control enforcement, consider a scenario where the patients assigned to Dr. Juny are Neil and Ryan. However, Dr. Juny is trying to access images of patient Mary. When the Action Agent discovers this mismatch, the access request of Dr. Juny is denied. Action Agent compares the access control policy rules with the user's Action Tuple and sends an access grant/denial response to the Authorization Server. If both responses from Patient Agent and Action Agent are positive, AS grants an Access Token to the user.

4.2 Behavior control

The Behaviour Agent in Figure 4 extracts the behaviour pattern of the user by analysing user attribute values that are stored in the Action Repository. By investigating the values of a particular attribute in the action tuple, the agent extracts the

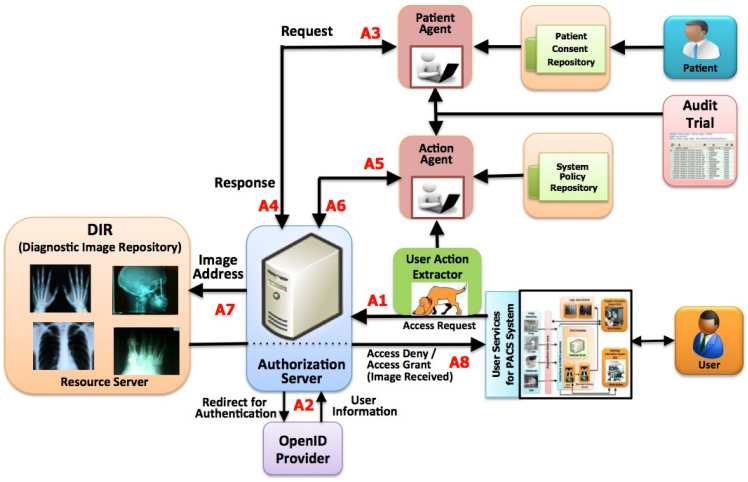

Figure 5: Access Control process using OAuth authorisation protocol.

behaviour pattern of the user. After deriving the behaviour of the user, the Behaviour Agent performs suitable pattern analysis to know the relevancy of the identified user behaviour. If the behavior is justifiable the agent sends appropriate messages to the security administrator who will modify the system policy rules accordingly.

Steps A to F in the proposed framework of Figure 4 correspond to updating access control policies based on user behavior. These steps are explained below. The Behavior Agent is responsible for: *Task1*: extracting user behavior patterns; and *Task 2*: selecting the significant behavior pattern to send to the System Administrator for updating the system security policies.

Step A. Every time the PACS user initiates a transaction for image retrieval, the "Access Control" component is notified to accumulate the corresponding action tuple instances for that session and send them to the "Action Repository" component for storage.

Steps B & C. The behavior agent extracts a large number of user's attribute sequence-patterns from the Action Repository (namely "behavior patterns"). Each behavior pattern is a frequent sequence of values that a particular attribute in the action tuple can hold. The extracted behavior patterns are stored in the "Behavior Repository" for further analysis. The behavior extraction procedure explained in Subsection 5.2.

Step D. The behaviour agent analyzes the user's behaviour patterns in order to identify the "significant behavior patterns" based on the size and frequency of the patterns.

Steps E & F. The behavior agent sends the significant behavior patterns to the system Security Administrator for monitoring the user activities and updating the system access control policies.

5 Case Studies

In this section, we describe two case studies which examine our proposed access control techniques discussed in Section 4. The access control technique is based on patient consent and user operation. However, the behavior-based sequence pattern extraction technique identifies common behavior of the users which will serve as a decision support mechanism for the system administrator to update the system's access control policies.

5.1 Action-based access control

As described in Section 3, OAuth authorisation begins with access grant approval followed by providing Access Token to the user services. In our proposed architecture, an access grant is given only after consulting with the Patient Agent and Action Agent (Subsection 4.1). In this section, we pro-

Table 1: Access control policies

	System Security Policies
1	Physicians are allowed to access to all kinds of images from the DI-r.
2	Only physicians can delegate nurses to look into certain medical images of patients as a part of diagnosis.
3	Specialist doctors can only look into the images of patients within their respective department otherwise delegated by a physician.
4	Nurses can only access the images of those patients that are assigned to them.
5	Lab technician has no right to access medical images of patients.
6	Medical Students can only access the medical images during their case study when delegated by their supervising physician.
7	Unless authenticated with the common infrastructure medical person is not allowed to access the images.
8	In some exceptional cases, user's access patterns can become a part of system security policies.
9	Users are allowed to access the images only when they are working in a medical environment.
10	Users will have to register their working hours during their specified shift.
11	Users have to specify the reason for accessing images, e.g., checkup, diagnosis, update, discharge, studies.

vide certain scenarios of access request by different medical professionals to access medical images in the DI-r. In this case study, we mainly consider the daily working patterns in a multi-specialty hospital and investigate the nature of image access by medical professionals, such as physicians, nurses, radiologists, medical students, lab technicians, cardiologists, gynecologists, physiotherapists, etc. Table 1 presents the access control policies that are defined by the Security Administrator for medical image retrieval from the DI-r.

Let us consider an example that includes a specific scenario of an access request made for image retrieval during a working day. Regional Medical Centre (RMC) is a multi-specialty hospital. The followings are the medical staff involved during a particular day. *Physicians*: Eric, Aadi; *Nurses*: Rona, Cole, Nims, Juny; *Medical students*: Jini, Ann, Aish; *Orthopedic surgeon*: Sherry; and *Patient*: Mike. *Locations from which images are allowed to access*: Hospital Information System (HIS), EMR Server, Radiology Work Station (RWS), Nurses stations (NS).

A portion of patient consent directives defined by patient Mike is provided in Table 2.

Table 3 describes a portion of access requests made by different users during a day. Table 4 presents how authorisation is granted by comparing: user access request, patient consents, and system security policies.

As the table indicates, access is granted only if both the patient consent is positive and the relevant system policy is satisfied. If consent is not defined for a particular case, we imply a positive consent and we only consider the system policies to make the access decision. On the other hand if the user access request matches with a system policy, but the patient consent does not approves that particular access request, the access request is denied. This is a simplified view of making a (granting/denying) decision on a particular user access request. There are more attributes in the Action Tuple that we extract when the user submits his access requests to the system. This includes user information, user role, user location, server location, time of day, team, delegation, requested profile status, service invocation type, etc. However, in this case study we have only included a few attributes.

Table 2: Example PCD of a patient.

Patient Consent Directive
a. Medical images should only be used for diagnosis.
b. Only designated nurses are permitted to access the medical images.
c. Nurses should not delegate access rights to other nurses.
d. Physicians can delegate access rights to other physicians.
e. Surgeons can access the images for diagnosis purposes.
f. Only care providers should access the images.

Table 3: Example access requests made by different users during a day.

	Access Request
I	Physician Eric logins to his account at a conference and tries to search MRI images of patient Mike.
II	Nurse Rona logins and tries to review CT scan images of patient Mike who is assigned to nurse Cole.
III	Nurse Cole delegates her authority to another nurse Juny to update her patient Mike's MRI Images.
IV	Orthopedic surgeon Sherry accesses X-ray images of Mike for diagnostic (delegated by Physician Eric).
V	Physician Aadi delegates nurse Nims (for patient discharge) to see fluoroscopy images of patient Mike.
VI	Physician Eric delegates his medical students Jini, Ann and Aish for a particular session to review the MRI images of patient Mike as a part of their case study.

Table 4: Authorization decisions made by comparing different parameters.

Access request No.	**Patient consent No.**	**System policy No.**	**Access Decision**
I	Not defined	9	**Denied**
II	b	4	**Denied**
III	c	2	**Denied**
IV	e	3	**Granted**
V	d	2	**Granted**
VI	a	6	**Denied**

5.2 Updating system policies using user behavior

Action tuple contains attributes that can take a particular value. In Figure 4 every time the PACS user sends an access request to the Access Control component, the attributes of the access request message are collected by the User Action Extractor module, and a new Action Tuple instance is created and stored in the Action Repository. At specific time intervals (e.g., a day, a week, or a month), the Behavior Control Agent retrieves from the Action Repository the sequences of attribute values for a single attribute in the Action Tuple. Then by applying a sequence pattern mining algorithm on those sequences the Agent extracts the behavior patterns of the different PACS users. We adopted the algorithm presented in a previous work of the authors [27] to extract behavior patterns. The sequence pattern mining algorithm produces a large number of patterns with different sizes and different frequencies which will overwhelm any useful analysis. Therefore, we should filter the patterns to identify the significant behavior patterns with larger sizes and higher frequencies. For this purpose, we set a threshold value to filter both insignificant frequencies and small behavior patterns. For example, if we set the frequency threshold value to 4 the Behavior Agent will keep the user behavior patterns with sizes 4 and above, meaning the behaviors have occurred more than four times during a specific time duration. Further, the Behavior Agent checks the extracted significant user behavior patterns against the rules in the System Policy Repository for situations such as: i) possible threats to the system information integrity; ii) possible improvement to some existing rules; and iii) lack of any rule to regulate the identified behaviors which requires adding new rules. If the Behavior Agent identifies one or more of the above situations, it reports to the System Security Administrator about the case, where the Security Administrator will be expected to investigate to update the security policy rules.

Such an approach will ensure a continuous and adaptable policy enhancement process based on the users actions and behaviors. It also reduces redundant access denials and improves the efficiency of the system transactions by reducing the number of requests and responses, while make it easier to detect any malicious or destructive system usage. In the following we define several user actions of nurse Rona from Radiology Department for six days and investigate the frequent action sequences. Table 5 presents the assigned tasks to nurse Rona.

Table 6 presents the list of actions done by Rona in six days at hospital. Extracting the tuples $< action, frequency >$ indicates the more frequent actions. These tuples in six days are:
$\{< A1, 6 >; < A2, 3 >; < A3, 1 >; < A4, 3 >;$
$< A5, 5 >; < A6, 4 >; < A7, 5 >; < A8, 3 >\}$

Considering a threshold value of 6, action A1 (occurred in all 6 days) satisfies the threshold value.

Consequently, the Behavior Control Agent eval-

Table 5: Several actions of a nurse.

Action
A1. Rona accesses ultrasound images of patients in the Obstetrics and gynaecology section.
A2. Rona is delegated by physician Eric to manage patients in Rheumatology section.
A3. Rona assists orthopedic surgeon Sherry in surgery.
A4. Rona is a radiology expert and leads the practical session for nursing students.
A5. Rona deals with patient discharge in all depts.
A6. Rona reviews and updates metadata of images
A7. Rona attends accident and emergency services.
A8. Rona has weekly duty in Neurology and MRI.

Table 6: Action sequence of nurse Rona.

	Action sequence
Day 1	A1, A7, A5, A6, A8
Day 2	A5, A1, A2, A7, A4
Day 3	A6, A7, A8, A1, A5
Day 4	A2, A4, A1, A7, A6
Day 5	A3, A5, A6, A2, A1
Day 6	A1, A7, A4, A8, A5

uates action "A1" as a justifiable action to be generalized as a common action and suggests to the administrator to consider it for updating the security policy of the users. The resulting policy may be as: "*Nurses can work with other departments if their qualification and practical experience is exceptional*". This policy loosens policy 3 and 4 in Table 1 and strongly supports policy 8 that makes the system flexible by providing a facility for dynamic updating of system policies by considering user's frequent actions.

5.3 Extracting behavior patterns

Action tuple contains attributes that can take a particular value. In order to extract the behavior of the user, we obtain the common sequence of values in an attribute. We use an algorithm from [27] to do this task. In our scenario, we consider the action of nurse (role) Rona (user) in RMC hospital (user location). She accesses the MRI images (requested data type) of patient Smith (requested profile status) by logging into the PACS in the radiology department (Server location). These attributes are fixed. The only attribute that is changing in the action tuple is "time of day". We analyze the access pattern of nurse Rona for different times within 7 consecutive days. Finally, we represent the behavior of Rona as a set of combinations of time attribute values with varying size (length) and the number of their occurrence. Let t_i represent various time of the day and i can take values from 0 to 23. For example, t_5 represents 5am, and t_{20} represents 8pm.

Table 7: Action sequences of nurse Rona.

Day	Time of Day
1	t5, **t7**,t10, **t12**, **t15**, t18, **t20**, t21
2	t1, **t7**, t8, **t12**, t14, **t15**, t19, **t20**
3	t4, **t7**, **t12**, **t15**, **t20**, t22
4	t0, t5, t10, t11, t18, t19
5	t2, **t7**, **t12**, **t15**, t16, **t20**
6	t1, t3, t8, t9, t13
7	t3, **t7**, t8, **t12**, 13, **t15**, t17, **t20**

Table 8: Behavior pattern of nurse Rona.

Size	Occur	Frequent Sequence
4	5	{**t7, t12, t15, t20**}
3	2	{t3, t8, t13} {t5, t10, t18}

Table 7 shows the time sequences during which nurse Rona accessed the medical images for seven continuous days. The length of sequence (number of elements in the sequence) and the number of their occurrences have linear relationship. We only consider those sequences whose length and occurrence values are greater than or equal to 3 and 2, respectively. Anything below these thresholds are considered insignificant. From the table we obtained the most significant combination that is t7, t12, t15, t20. It means that Rona accessed the images at 7am, 12pm, 3pm and 8pm for five of the seven days of observation. We now use this interaction pattern to derive her behavior. Later we investigate this behavior to justify her access nature. At the end of the analysis if the behavior is found to be significant, we will add this behavior as a part of system policies. By doing so, we can loosen constrains placed on the permission given to nurses to access the medical images of patients during var-

ious time interval of a day. On the other hand as security is concerned, if the behavior is found to be suspicious further investigation can be conducted. This makes our system dynamic in the sense that user behavior itself plays a role in modifying the access policies of the system.

6 Conclusion

This paper contributes to the domain of medical imaging by providing a solution for security and privacy aspects of the sharing of these images. The approach uses multi-agent systems which communicate through repositories and together provide an advanced mechanism for flexible and dynamic enhancement of system security policies. An action-based access control mechanism and an user-behavior based policy enhancement procedure have been proposed. The solution utilizes modern authentication and authorization techniques (OpenID and OAuth) that are applied through dedicated software agents. Patient consents are defined off-line by the patient and stored in the consent repository. The action agent makes the access control decisions by capturing the user operations. As the current PACS have closed architecture, capturing the user identification and requested operation are major challenges. Different network traffic analysis tools with filtering capabilities are required to collect such information to be used for the access control purposes. However, the availability of advanced techniques is a major driver for this project with obvious benefits in reducing huge costs of the existing PACS for safe communication and sharing of images with the DIr systems.

References

[1] Canada health infoway. https://www.infoway-inforoute.ca/index.php/resources/technical-documents/ [23 April 2013].

[2] Canada health infoway news. https://www.infoway-inforoute.ca/index.php/news-media/2011-news-releases/neodin-completes-northern-connections/ [23 April 2013].

[3] Ehr blueprint. https://www2.infoway-inforoute.ca/Documents/EHRS-Blueprint-v2-Exec-Overview.pdf/ [23 April 2013].

[4] EHRi privacy and security conceptual architecture. https://knowledge.infoway-inforoute.ca/EHRSRA/doc/EHR-Privacy-Security.pdf/ [23 April 2013].

[5] Gta west diagnostic imaging repository. http://www.gtawestdir.com/ [23 April 2013].

[6] Hospital diagnostic imaging repository services incorporated (hdirs). http://www.hdirs.ca/ [23 April 2013].

[7] (icam) openid 2.0 profile - idmanagement.gov. http://www.idmanagement.gov/documents/ICAM_OpenID20Profile.pdf/ [23 April 2013].

[8] Ihe xds-i. http://wiki.ihe.net/index.php?title=Cross-enterprise_Document_Sharing_for_Imaging/ [23 April 2013].

[9] The oauth 2.0 authorization protocol. http://tools.ietf.org/html/draft-ietf-oauth-v2-31// [23 April 2013].

[10] Southwestern ontario diagnostic imaging network. http://www.swodin.ca/ [23 April 2013].

[11] X. Chen et al. Hippa's compliant auditing system for medical imaging system. In *IEEE Conference on Engineering in Medicine and Biology Society*, pages 562–563, 2005.

[12] S Ferreira et al. Improving expressiveness of agents using openehr to retrieve multi-institutional health data: Feeding local repositories through hl7 based providers. In *IEEE Conference on Information Systems and Technologies, CISTI'12*, pages 1–5, 2012.

[13] Christos K. Georgiadis et al. Flexible team-based access control using contexts. In *ACM symposium on Access control models and technologies*, pages 21–27, 2001.

[14] S Gupta and S Pujari. A multi-agent system based scheme for healthcare and medical diagnosis system. In *IEEE Conference on Intelligent Agent & Multi-Agent Systems, IAMA'09*, pages 1–3, 2009.

[15] P. Jayaratna and K. Sartipi. Hl7 v3 message extraction using semantic web techniques. *International Journal of Knowledge Engineering and Data Mining*, 2(1):89–115, 2013.

[16] Jian Jiang et al. A federated identity management system with centralized trust and unified single sign-on. In *IEEE Conference on Communications and Networking in China, CHINACOM'11*, pages 785–789, 2011.

[17] Campos Maria João, Manuel E. Correia, and L. Antunes. Leveraging identity management interoperability in ehealth. In *IEEE Conference on Security Technology, ICCST'11*, pages 1–8, 2011.

[18] Peyton Liam et al. Addressing privacy in a federated identity management network for ehealth. In *IEEE World Congress on the Management of eBusiness, WCMeB'07*, pages 12–12, 2007.

[19] Z Lixin. Active health evaluation with multi-agent. In *IEEE Conference on Genetic and Evolutionary Computing, WGEC'09*, pages 169–172, 2009.

[20] M. Mario, K. Delac, and M. Grgic. Overview of the dicom standard. In *IEEE 50th International Symposium on ELMAR'08. Vol. 1*, pages 39–44, 2008.

[21] Hecht Maximilian. PACS-Picture Archiving and Communication System. Master's thesis, Vienna University of Technology, University of Paderborn, Austria, 2008.

[22] Deng Mina et al. Identity in federated electronic healthcare. In *IEEE Wireless Days, 2008*, pages 1–5, 2008.

[23] S Rossilawati et al. E-health services with secure mobile agent. In *IEEE Conference on Communication Networks and Services, CNSR'09*, pages 270–277, 2009.

[24] J Tian and H Tianfield. Health delivery systems—a case for multi-agent systems. In *IEEE Conference on Systems, Man and Cybernetics, SMC'09*, pages 2718–2722, 2009.

[25] Sofia K Tzelepi et al. A flexible content and context-based access control model for multimedia medical image database systems. In *ACM workshop on Multimedia and security: new challenges.*, pages 52–55, 2001.

[26] M.H. Yarmand, K. Sartipi, and D.G. Down. Behavior-based access control for distributed healthcare environment. In *Symposium on Computer-Based Medical Systems*, pages 126 – 131, 2008.

[27] Mohammad H. Yarmand, Kamran Sartipi, and Douglas G. Down. Behavior-based access control for distributed healthcare systems. *Journal of Computer Security*, 2(1):1–39, 2013.

[28] Stephen S. Yau et al. Situation-aware access control for service-oriented autonomous decentralized systems. In *IEEE Conference on Autonomous Decentralized Systems, ISADS'05*, pages 17–24, 2005.