RJ10382 (A0606-004) June 6, 2006
Computer Science

IBM Research Report

The Eclipse Open Health Framework

Srivatsava R. Ganta, Eishay Smith, Sarah E. Knoop,
Sondra R. Renly, James H. Kaufman
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

The Eclipse Open Health Framework

Srivatsava R. Ganta, Eishay Smith, Sarah E. Knoop, Sondra R. Renly, James H. Kaufman
IBM Almaden Research Center
650 Harry Road, San Jose, CA 95120

Abstract

The success of open source software in other fields has generated a surge of interest in open source for healthcare. Open source initiatives not only help reduce costs for the patients and healthcare providers but also establish standards that are necessary to prevent vendor lock-in. The National Health Information Network (NHIN) project is aimed at modernizing the US Healthcare Information Infrastructure (HII) by integrating various regional health centers. Such an initiative requires standardization of protocols and data models. Eclipse Open Health Framework (OHF) is the official open source organization for HL7 and one of the leading efforts to deliver an open source, standards-based platform, for healthcare software. In this paper, we present an overview of the OHF project, describe the technologies it will provide in the near future, and discuss possible future directions for this important community.

1. Introduction

With the rise of the internet and its enabling of diverse production models and communication paths, open source software (OSS) development is becoming strategic throughout the IT industry. The open source phenomenon has not only affected the software industry but also many other fields of science and technology such as biotechnology [7][8][9]. Open source software development is driven by a community of developers with shared interests. The goal is to reach a consensus on the community requirements and develop software that can be made publicly available. These requirements are mutually agreed upon standards that are vendor-neutral. The term open source is in general understood to imply as software whose source code is open. However, there is much more to it than that:

1. Software Licensing Models – Open source software comes with a variety of licenses. Most grant absolute copyright to the user though the terms and conditions vary.
2. Software Development Model –The open source development model differs from the traditional software development and places emphasis on collaboration among the community members.
3. Business Model – The open source software development enables emerging business models based on revenue generated from support services, consultancy services and product enhancements.

Open source efforts in healthcare started as early as 1980 [20], however, web-based projects began around 1998 during the height of dot-com period. Several websites offered open source Electronic Medical Record Systems (EMRs), Physician Practice Management (PPM) packages, Diagnostic Image-Viewers etc. These initiatives slowly became dysfunctional owing to the dot-com bubble burst, minimal or poor implementations of standards, and lack of interest from the healthcare community. The OpenEMR[6] project was then started by the American Association of Family

Physicians, but did not attract a developer community. Recently, the proven success of open source in other fields, and the government initiative to modernize Healthcare Information Infrastructure (HII) [3], lead to a renewed surge of interest in open source for healthcare. Although the healthcare industry was late in adopting open source, several projects such as OpenEMR[4], OpenVista[5], OpenEMED[6], OpenMRS[16], EGADSS[17], Oscar[18], PICNIC[19] are currently underway and are expected to have a huge impact on the current healthcare system. Government agencies along with dominant healthcare enterprises are looking to open source software as a vehicle to reduce costs for patients and healthcare providers.

The U.S. healthcare industry is by far the largest in the world in both absolute dollars and in percentage of GDP (more than $1.5 trillion, or 15 percent of GDP). It is also a highly fragmented and complex industry with numerous providers, insurance agencies and other players. This scenario, coupled with an antiquated infrastructure for the collection of and access to medical data, lead to enormous inefficiencies and sources of error over the decades. Recently, there has been increased focus on the need to modernize the healthcare information infrastructure in the United States. Consumer, regulatory, and governmental pressure drive a growing consensus that the time has come to modernize the U.S. healthcare information infrastructure (HII). The National Health Information Network (NHIN) [15] is a project undertaken to integrate the information infrastructure of the various healthcare providers in the country. It is expected to transform the way healthcare technology works. This transformation is envisioned to improve healthcare quality, expediency, efficiency and successful delivery of healthcare while decreasing costs.

The Open Health Framework (OHF) [21] is an open source initiative focused on developing standards-based platform towards the integration goals of NHIN. The project addresses the need to improve levels of interoperability between applications and systems within and among healthcare organizations. More specifically, it provides implementations of key health informatics standards based components in a plug-in based model. The overall goal of the framework is to provide a platform which could become the core foundation for a multitude of healthcare application developers. In this paper, we present an overview of the current state of OHF and discuss some future directions. The rest of the paper is organized as follows: Section 2 gives an overview of the current state of healthcare IT market; Section 3 presents the barriers to the realization of a NHIN. Section 4 discusses the current state of OHF at the conceptual level and some possible future directions. Section 5 provides the conclusion.

2. HealthCare IT Market Scenario

Over the decades, the Healthcare IT (HIT) market has been dominated by few large enterprises and has lacked competition due to vendor lock-in and proprietary systems. Consequently, there has been little opening for the market to grow and incubate innovation. A mature market can be defined as the one that reached a state of equilibrium marked by the absence of significant growth. The HIT market, though existing for a while, is far from being mature and lacks competition and hence affordable products. This resulted in the delayed emergence of technology in the US Small Medium Clinics (SMCs) market. This is evident by the fact that only 12% of SMCs have EMR systems, while the rest still use antiquated paper forms and filing systems [25].

The two predominant software products available in the HIT market are: 1. Server End products (SEP), such as PACS Image Archive; 2. Information Consumer and Producer End (ICP) products, such as embedded software in medical devices, EMRs etc. Although the trend is changing, HIT enterprises selling SEP products tend to provide holistic solutions which bundle tuned ICP and SEP systems. Such single-vendor systems stimulate the use of proprietary interfaces between the SEP and ICP components in order to achieve better integration and performance which gives competitive advantage to large enterprises. The incentive to open up the system by supplying a full interface and proprietary data model is small, and supporting a standard interface is even smaller for these organizations. This makes it difficult to replace subcomponents of such a system and leads to vendor lock-in, which raises the business and technical barriers for new competitors to enter the market segment. For instance, in the context of radiology systems, Picture Archiving Communication Systems (PACS) have been available in the HIT market from some time. Most PACS vendors provide a full set of products including imaging machines (modules), display machine, image archives, and other satellite components. The vendors create a tight integration between the software and hardware. Though the Digital Imaging and Communication in Medicine (DICOM) standard is well defined, and is not new to the market, the support of a full standard DICOM interface is not yet ubiquitous.

This scenario is analogous to the early days of computing when several non-standard operating systems, protocols, storage subsystems existed in the Information Technology market. There is a need for plug-n-play products that are available at affordable prices. The two main technological barriers to migrating to a plug-n-play environment are:

1. Product complexity: Since the market is not new, and the product requirements are very complex, reaching a good competition level of products requires years of development.

2. Interfaces to backend: Large companies which dominate healthcare SEP products have little motivation to supply full standard interfaces.

So far, the financial center of gravity in the HIT market was in the large healthcare enterprises. However, analyst reports [1] [4] show the dynamics have shifted and that today, in the U.S., SMCs collectively spend more on HIT than large enterprises. This calls for a much needed paradigm shift in the market, and an initiative to encourage open standards that help ensure competition and innovation.

3. Towards a National Health Information Network (NHIN)

The concept of NHIN has evolved over several decades with a goal of providing an electronic health record for every American within the next decade. The idea is to link disparate health care information systems together to allow patients, physicians, hospitals, public health agencies and other authorized users across the nation to share clinical information in real-time under stringent security, privacy and other protections. While such transformation may be disruptive in the short term, it will, in the future, significantly improve the quality, expediency, efficiency, and successful delivery of healthcare while decreasing costs to patients and payers and improving the overall experiences of consumers and providers

Healthcare delivery is executed using various economic models in different countries around the world. Centralized markets involve a single healthcare provider, most often the government which runs the information infrastructure for the whole nation. For example,

countries like Brazil have a single point controlled healthcare system. Semi-centralized markets involve a close group of contractors who work together with the government guidelines to achieve a common goal. This scenario could be seen in countries such as the UK. However, in the US, the healthcare industry is heavily decentralized and operates based on independent solutions offered by disparate enterprises. This scenario poses a greater challenge in the realization of a NHIN than when compared with the other two market scenarios as it would involve achieving interoperability among highly heterogeneous systems.

According to the IEEE Standards Computer Dictionary, interoperability is the "ability of two or more systems or components to exchange information and to use the information that has been exchanged". Creating an interoperable NHIN requires a healthcare software ecosystem in which it is relatively easy for any Independent Software Vendor (ISV) product to produce, acquire, and transfer medical information. To achieve such an environment in the HIT industry, two key things need to happen: First, the standardization of healthcare information-transfer protocols and healthcare document structures; Second, the emergence and adoption of publicly available healthcare software tools.

3.1. Standardization

To enable the entry of new ISVs and incubate innovation in a market, low level standards need to be present. A good example for this would be the emergence of internet based on standardized interfaces. Lack of low-level standards leads to proprietary solutions and thereby vendor lock-in. Further, the current state of the HIT industry makes this task tougher because of several reasons:

• Product Legacy: The usage of products which have been in existence for some time and predate the emerging standards makes migration an expensive task. This would mean that a lot of proprietary code (and perhaps functionality) will render useless and adjusting to the new workflow and data model standards might need a lot of effort.

• Chicken and Egg Syndrome: Most EMR and PHR system providers will not support public standards before having a solid business case. In other words, they would not be willing to adopt public standards before their clients have SEP products supporting those standards. This is more true to small companies which do not have a large budget to cover harmonization efforts. To be in the market, an ISV would have to add support for the specific interface exposed by the SEP products used by their clients whether it is a standard or not. This in turn results in increased product costs because of the extra burden that goes into maintaining and adding these interfaces for every SEP product in the market. On the other hand, the large number of ICP products depending on its proprietary interface lowers the motivation for an SEP product provider to shift to standard interfaces.

• Rapid Change in Standardization: Most healthcare standards are not yet mature and have been changing at a relatively high pace. This makes it extremely for a small ISV to keep in pace with evolving standards which changes once in few months, or support several different version of the same standard.

• Lack of Comprehensive Standards: Healthcare standards, like most standards, may not fully cover a product’s requirement. In such cases, the providers use proprietary protocols in addition to the standards which again leads to vendor lock-in. An example of this can be taken from the internet browser market. The addition of non-standard web

components such as ActiveX led to scenario in which only browsers which could handle the non-standard add-ons could render the page.

• Added Effort towards Implementation: Most of the ICP products need to use only a subset of a standard. However, because of inherent dependencies, the product provider would need to implement the whole standard which takes up a lot more man-power they find they need to implement a lot more than is required for the product due to lack of standard libraries.

3.2. Open Source Tools

Development of tools that enable software interoperability among heterogeneous systems generally involves implementation of standardized interfaces. This could be given a boost through the availability of publicly available libraries that implement the tools. For example, in the web service development environment, availability of open source libraries that implement the web service standards boosts the number of ISVs that venture into the development of products that use these libraries.

Access to such publicly available healthcare standard-based tools is critical for the establishment of an NHIN in a decentralized market. These libraries could provide software that implements standard healthcare document architectures, message protocols etc. These tools could free ISVs from integration efforts so they can focus instead on implementation of their business and application logic. As mentioned earlier, HIT standards change at a rapid pace. Maintaining a standards-based product hence becomes prohibitively costly for small ISVs. Availability of up-to-date third party tools would help them to keep up with the current standards. Finally, this will eliminate vendor lock-in and allow ISVs to compete and innovate.

4. Eclipse Open Health Framework (OHF)

The goal of Eclipse Open Healthcare Framework (OHF) [21] project is to create an open-source platform that provides standards-based software to improve levels of interoperability between various healthcare enterprises. The aim is to provide a platform which could become the core foundation for a multitude of healthcare application developers. The framework consists of a complementary set of tools and libraries that expose these functionalities as simple Java API and also in the form of Eclipse Rich Client Platform (RCP) plug-ins [22]. The Eclipse RCP architecture facilitates design based on a plug-n-play model where new components can be created using the “extension points” of existing components.

The design of the OHF is based on “integration profiles” proposed by the Integrated Healthcare Enterprise (IHE) [23] organization, a multi-enterprise initiative designed to stimulate the integration of the healthcare information systems using existing standards such as HL7, DICOM etc. The framework identifies a set of functional components of the distributed healthcare environment called “Actors”, and specifies their interactions in terms of a set of coordinated, standards-based “Transactions”. These Actors and Transactions are abstractions of real-world healthcare information system environment. For example, if Organization A is requesting radiology documents from Organization B, A and B are the actors and the transaction would be “Request for Radiology Documents”. The rest of this section gives a technical overview of OHF components.

4.1. Cross-Enterprise Document Sharing (XDS)

The Cross-Enterprise Document Sharing (XDS) profile provides a standards-based mechanism to share clinical documents among a set of healthcare enterprises. These enterprises could range from a private physicians office to an acute care in-patient facility. The enterprises are assumed to be a part of a "Clinical Affinity Domain", a group of enterprises that have agreed to work together using a common set of policies and share a common infrastructure. An XDS document is a content neutral concept of a clinical document which can be plain text, formatted text (HL7 CDA etc) or an image (DICOM etc).

The profile functionality is based on ebXML Registry standards, SOAP, HTTP and SMTP. The document sharing is managed through federated document repositories and a document registry to create a longitudinal record of information about a patient within a given clinical affinity domain. The functionalities are:

- A document repository is responsible for storing documents in a transparent, secure, reliable and persistent manner and responding to document retrieval requests.
- A document registry is responsible for storing information about those documents so that the documents of interest for the care of a patient may be easily found, selected and retrieved irrespective of the repository where they are actually stored.

Along with the Registry and Repository, the key actors in this profile are the "*Document Source*" and "*Document Consumer*". Various transactions between these actors define the functionality of the XDS component. Interested reader is referred to [24] for further details.

The current OHF package includes implementation of Document Source and Document Consumer actors along with the necessary document metadata processing module (XDS Metadata). The OHF XDS profile interacts with various security, authentication and auditing features from other components in the framework to provide a standards-based sharing of clinical documents.

4.2. Patient Identification

In the US there is an intentional lack of a national identification system. Therefore, significant effort in healthcare goes into ensuring that each encounter is correctly associated with the correct patient's medical record. In some cases this leads to multiple medical record numbers for one individual when insufficient data is available. Also, an individual seeking multiple services, even within one organization, can result in many patient record identifiers due to the operation of different systems.

The OHF Patient Identifier Cross-referencing (PIX) and the Patient Demographics Query (PDQ) Profiles aim to support the ubiquitous number of patient identifiers through standardized transmission of patient identity information into a repository and reciprocal query of that information to return the appropriate domain specific medical record number. This is accomplished through three profile transactions. The repository input, or PIX Feed, is accomplished using existing HL7 standard ADT message types for registering, admitting and updating a patient record. The reciprocal query transaction, or PIX Query, uses a HL7 standard query/response message for getting and receiving the domain corresponding patient identifier. A more generic query transaction, or PDQ Query, also uses HL7 query/response message pairing to request

patient identifiers matching a set of demographic information. For example, return patient identifiers for patients that have lived at a particular address.

The current state of the OHF framework includes the PIX Feed, PIX Query, and PDQ Query client-side transactional support in the form of plug-ins that support a variety of data input options to provide for the range of system capabilities existing today. In adopting the HL7v2 message object, these clients have the added capability of running a conformance protocol verification step. Any problems detected are documented in a log and the message can be withheld from submission depending on the level of error generated. The implementation is operating with MLLP (Minimum Lower Level Protocol) for client to server communication and is open to support additional protocols in the future.

4.3 Other Components

In addition to patient identification and document exchange, other key functions necessary for healthcare interoperability are messaging standards, access to terminology services and secure user authentication. The OHF framework provides the implementation of HL7v2 and the newly released HL7v3 message standards to facilitate healthcare applications and other integration components with an up-to-date interface to create and manipulate HL7 messages.

Accessing and utilizing terminology resources are a common and necessary function for many healthcare IT applications ranging from decision support to EHR integration. It is, therefore, an important feature needed for interoperability within and across healthcare enterprises. The challenge of defining a predictable set of vocabulary APIs has been addressed by the HL7 Common Terminology Services (CTS) standard. The CTS defines the minimum set of functions required for terminology interoperability within the scope of HL7's messaging and vocabulary browsing requirements. The OHF framework provides an API implementation for CTS and thereby enables query and access of terminological content.

Many of the current healthcare IT applications use different technologies and platforms that include Microsoft .NET, MUMPS and LAMP(Linux, Apache, MySQL, PHP). OHF provides a Web-Services based interface through the OHFBridge component which allows existing applications to access the OHF interoperability profiles. Secure user authentication across enterprises is provided by the implementation of IHE (Enterprise User Authentication) XUA integration profile.

4.4 Ongoing Work

Privacy is one of the foremost concerns for healthcare organizations towards the participation in an integrated system. In the context of an IHE based scenario, privacy concerns arise from the two main actors: *Document Source* and *Document Consumer*. The *Source* is concerned with addressing the privacy policy of its patients when sharing data with other enterprises. This would include reservations on sharing a set of documents or a subset of the document with other users. Though the problem of granting access to eligible users can be solved through authentication, the issue of applying privacy preservation policies still remains. On the other hand, a more subtle, yet important privacy concern exists for the *Consumer*. This includes the breach of privacy that could occur if one traces the way consumer accesses the shared data. For

example, if one can trace that a *Consumer* is trying to access CD4 test results of a patient, it could be inferred that the patient is infected with HIV.

The OHF community is working on addressing this problem by associating the "privacy policy" of each actor (*Source* and the *Consumer*) with the access to the shared documents in the system. This can be achieved using the emerging industry standard of Platform for Privacy Preference (P3P) [26], an initiative from the W3C [27] standards organization. This would enable the actors to define their corresponding privacy policy which can be associated with all or some of the documents shared by that actor. The system functionality then imposes the requirement of policy "compatibility" for any access or querying to happen. This component is aimed to provide the basic privacy preservation functionalities for an integrated healthcare system.

5. Conclusion

The success of Open source software in various fields is bringing a new surge of interest in open source efforts for healthcare technology. The realization of an Integrated Healthcare System requires a market scenario with open standards and publicly available tools. The Eclipse Open Health Framework (OHF) is the official open source organization for HL7 and one of the leading efforts to deliver an open source, standards-based platform, for healthcare software. The project provides the basic tools and libraries needed for the development of a healthcare application in an interoperable environment. The vision is to bring a community into existence that would work towards the goals of a fully integrated healthcare system.

6. References

[1] Eclipse Open Health Framework Platform. http://www.eclipse.org/ohf. Last visited March 2006.

[2] Common Public License. http://www.eclipse.org/legal/cp1-v10.html. Last Visited March 2006.

[3] http://www.whitehouse.gov/infocus/healthcare/. Last visited March 2006.

[4] http://www.openemed.org/. Last visited May 2006.

[5] http://www.oemr.org/. Last visited May 2006.

[6] http://sourceforge.net/projects/openvista. Last visited May 2006.

[7] http://www.cambia.org.au/daisy/cambia/home.html. Last visited May 2006

[8] http://creativecommons.org/. Last visited May 2006.

[9] http://en.wikipedia.org/wiki/Wikipedia. Last visited May 2006.

[10] Lerner, Josh and Tirole, Jean, "Some Simple Economics of Open Source" . Journal of Industrial Economics, Vol. 50, No. 2, pp. 197-234, June 2002 Available at SSRN: http://ssrn.com/abstract=313493

[11] http://www.informatics-review.com/open.html. Last visited May 2006

[12] http://www.ieeeusa.org/policy/positions/NHIInteroperability.html

[13] http://www.whitehouse.gov/infocus/healthcare/ Last visited March 2006

[14] US Physicians:Technology Usage and Trends. eMarketer January 2005

[15] A Strategy for Building the National Health Information Infrastructure. http://aspe.hhs.gov/sp/NHII/Documents/NHIIReport2001/default.htm ; Last visited March 2006

[16] http://openmrs.org/wiki/Main_Page

[17] http://egadss.sourceforge.net/

[18] http://www.oscarmcmaster.org/

[19] http://picnic.euspirit.org/

[20] http://ehr.gplmedicine.org/index.php/FreeMED_Review

[21] http://www.eclipse.org/ohf/

[22] http://wiki.eclipse.org/index.php/Rich_Client_Platform

[23] http://www.ihe.net/

[24] http://www.ihe.net/Technical_Framework/upload/ihe_iti_tf_2.0_vol1_FT_2005-08-15.pdf

[25] Mark Leavitt (Certification Commission for Health Information Technology Chairman) @ AHCJ 2006

[26] http://www.w3.org/P3P/

[27] http://www.w3.org