A Lightweight Approach to Interoperable Health Information Exchange

Anthony W. Wallace and Rex E. Gantenbein
Center for Rural Health Research and Education
1000 E. University Avenue, Department 3432
University of Wyoming, Laramie WY 82071 USA
E-mail: wallace@uwyo.edu, rex@uwyo.edu

Abstract

The ability to electronically exchange health information among different facilities or organizations has significant potential to improve patient care and reduce the costs of that care. However, the only approaches currently known to support this exchange of information either rely on expensive interface engines or are dependent on a centralized repository of data. We propose instead to use existing Web technology to develop a lightweight and relatively inexpensive mechanism for exchanging information that allows organizations to maintain their records independently but supports requests from other organizations for information to be fulfilled.

Keywords: electronic health records, information exchange, system design, distributed systems, databases

1.0. Introduction

In April 2004, U.S. President George W. Bush issued an executive order establishing the Office of the National Coordinator for Health Information Technology. This office was created to support developing a nationwide interoperable health information technology (HIT) infrastructure. It is widely held that the development of HIT has the potential to greatly improve the quality of patient care, enhance the efficiency of health care professionals, and reduce the costs of health care. At the same time, the U.S. health care system has been slow in its adoption of HIT due to a number of challenges, including cost, resistance to change, concern of security and privacy, and unfamiliarity with the underlying technology.

An important factor in encouraging widespread adoption of HIT is the ability for providers to exchange information among varying forms of electronic health records (EHRs) and health information management systems. The U.S. government has issued directives mandating that health care providers must have systems capable of exchanging data with each other by 2014 [1]. However, finding a way in which the disparate systems in existence today can exchange information with each other has the potential to be both timely and costly. Most currently proposed solutions require a connection from each system to every other system and/or the use of interface engines to achieve such exchange.

As an alternative, we propose a lightweight solution for health information exchange based on Web services technology. Our solution is flexible, relatively affordable, and can operate over the commodity Internet. Our solution is also scalable to a large number of systems and is able to deal with changing communication standards and data sources.

2.0. Interoperability and Electronic Health Records

There are a number of different views as to what comprises an electronic health record. An Australian EHR taskforce defined the EHR [2] as:

> *an electronic longitudinal collection of personal health information usually based on the individual, entered or accepted by health care providers, which can be distributed over a number of sites or aggregated at a particular source. The information is organized primarily to support continuing, efficient, and quality health care. The record is under the control of the consumer and is stored and transmitted securely*.

A number of EHR systems are already commercially available throughout the U.S. and other countries; however, these systems are largely based on proprietary architectures and are largely incompatible with each other in terms of sharing or communicating information. Standards are being developed to define a common framework for exchanging information, but as yet there is no equivalent of internetworking for health information systems. This lack of *interoperability* eliminates one of the major advantages of using EHRs, the ability to quickly and accurately transmit a given individual's information among various providers and facilities working with that individual. It also prevents the individual patient from having easy access to all of his/her health information, an idea that is gaining support as a way to provide patient-centered health care and involve the patient in the decision-making process throughout the continuum of care.

There are two basic approaches to supporting interoperability for EHR systems. The *centralized* approach is based on collecting and storing comprehensive patient data in a single repository or location. The most extreme form of this approach is a single data repository containing the health information for every person that ever saw a practitioner. While easy to comprehend, such a system would have to be implemented at a national level (if not global); not only would this be a massive undertaking to implement and maintain, it would likely meet significant resistance on the part of practitioners and patients alike.

A more reasonable approach could be to use some kind of interface to integrate subsets of information. Here, individual service providers maintain the details of their encounters with the patients in independent records, which are then integrated into a comprehensive EHR. Authorized individuals (including the patient) are able to locate and access these subsets of information through the interface and view the details of available records.

Unfortunately, in the modern health care system, there are likely to be a number of different providers involved in the care of a given patient. Depending on the complexity of the care and/or conditions involved, each of these providers may have an independent EHR for a single patient. This leads to the consideration of a *distributed* approach to the EHR, where no location is considered the "primary" repository of information. In this approach, a comprehensive record can be created only by locating the various records for a patient and integrating them at the location where the information is needed.

This approach has several advantages, including the ability to access the most recent entries in a patient's health history regardless of where those entries reside and the elimination of the need for a large centralized database, a tempting target for unauthorized users. However, it also requires the development of both technology and standards for exchanging information among the various information repositories.

Several attempts to develop health information exchange systems have been made over the last four decades, with varying results [1]. For example, the Mayo Clinic took over ten years to implement an electronic solution that supports a paperless office. This system, however, is still unable to exchange information between the clinic's facilities in Minnesota, Florida, and Arizona. The major problem in the lack of interoperability is that national standards for digitizing and interconnecting patient data have not been agreed upon. The Mayo Clinic, instead of creating its own standards of interoperability and running the risk of not being able to communicate with health care providers outside of its partnerships, has chosen to wait for national standards to be set before implementing an exchange mechanism.

Such concern about being able to communicate with health care providers outside of a single enterprise or organization is a valid one. A single cost-effective solution that can exchange information among multiple providers is clearly preferable over multiple interfaces designed specifically for individual health care providers. This would require standardizing the communications protocols used in exchanging information, eliminating the need for specific interfaces, but because standardization is focused on the communications, existing systems are not forced to conform to an internal or representational standard. Instead, internal information can be translated into the communication specification, so each participating facility can represent information as desired.

Standardizing the communication at the network level ensures that the information can be exchanged among health care providers; however, it does not guarantee that the information exchanged between providers represents the same thing. One approach to this problem is the HL7 protocol [3]. This protocol standardizes messages passed back and forth between entities within a health care enterprise as a string that contains several fields concatenated together, separated by a pipe delimiter [4,5].

Another approach to standardization is being taken by the Healthcare Information Technology Standards panel, which recently identified an *initial* set of 90 medical standards out of an list of about 600 [1]. These are the technical standards for the U.S. nationwide record system. However, the fact that this is defined as an initial set of standards suggests that these standards themselves may change.

3.0. Implementing Health Information Exchange

Even when standards are adopted, there will still be a need for an infrastructure through which standardized information can be exchanged. Interface engines have been at the heart of the exchange of medical information for many years [6]. The ability to automatically convert information between differing formats through a software interface is useful for information exchange. However, these systems typically are designed for interfaces to legacy databases within an organization and not for exchanging information across organizational boundaries. Moreover, they are very expensive to purchase and maintain, which puts the technology out of reach of small health care facilities, which comprise the major of practices in the U.S.

3.1. XML and Information Exchange

A less expensive alternative, in terms of both cost and development, is the use of Extensible Markup Language (XML) to provide Web-based interfaces to information. XML has been in development since 1996. Derived from Standard Generalized Markup Language, it was originally used to describe a class of data object, as well as to describe the behavior of programs. Since then, it has been found to be very useful in exchanging data. Using tags to surround a piece of data, XML documents are able to name objects within the document. This makes an XML document readable to both the human and the computer.

Several technologies accompany XML to support validating and reformatting XML Documents. Data Type Definitions (DTDs) are used to validate XML documents meet a specified structure. XML Style Language (XSL) applies formatting tags to an XML document so that the document can be displayed as desired. Parsers come in two varieties, Document Object Model, and Simple API for XML (SAX). To programmatically create XML documents, the Document Object Model (DOM) can be used, or a specific generator can be created.

An XML document exists as a file, and can be sent via FTP, HTTP, and e-mail. In addition, the Simple Object Access Protocol (SOAP) was created for transporting an XML documents via the network. A SOAP message contains three parts: an envelope, a header, and the body. The envelope is the basic container for the message. The header can contain information about routing the message, as well as any security measures that may need to be taken, while the body contains the XML message itself. SOAP messaging, which operates on the HTTP protocol, can be used to send information either in one direction, such as sending lab results to another department/process, or in two directions, such as a remote procedure call that requests and returns data.

3.2. Web Services and Information Exchange

Web services, first introduced by Hewlett Packard in 1999 and now available in a variety of development platforms (including Java and .NET), provide a simple way to transport information over the Internet from one platform/source to another platform/source using XML documents within SOAP messages [7]. Once deployed, Web services can be invoked remotely to retrieve or store information.

As shown in Figure 1, in order for an external application to use (consume) a Web service, the producer of the service provides a Web service description language (WSDL) file describing the service. The WSDL file contains XML code to describe the procedures that are available, what parameters are expected, the data types of the parameters, what the return values are, and the data types involved. The WSDL file is placed on a Web server where others may view it and access it to determine the service(s) available. Integrated Development Environments that can view the WSDL file over the Internet can automatically generate code to send and receive the messages as defined in the WSDL.

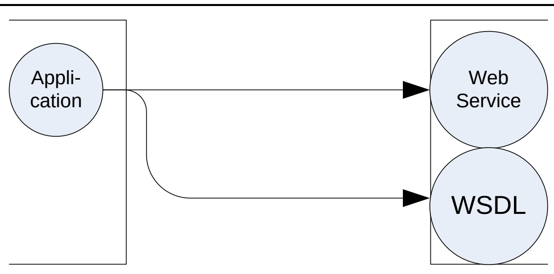

Figure 1. Basic consumption of a Web service by an application. The WSDL describes the interface of the Web service enabling the application to consume the methods offered by the Web service.

A basic Web services architecture places the Web server hosting the Web services outside the firewall. In addition to the Web service, database drivers are used to communicate with a database server located behind the firewall. However, this architecture is vulnerable to security breaches if the Web server is compromised, allowing the database server to be compromised as well.

Another approach has two Web servers, one outside the firewall, and one inside the firewall. The outside firewall hosts a Web service that is public to the world, and communicates with another Web service that is hosted on the Web server inside the firewall. The communication between these two Web services is done using SOAP and XML. The Web server inside the firewall enables the internal Web service to talk to the database driver. If the outside Web server is compromised using this architecture, the intruder still has no direct access to the database.

3.3. Interoperability and Web Services

Unfortunately, this approach still does not deal with the problem of interoperability across organizational or facility borders. To address this, we propose a lightweight and relatively simple solution to this problem using Web services, SOAP, and XML. The solution is also scalable, allowing the addition of new entities to the network with minimal impact to the existing entities. This solution allows health care providers participating in an exchange system to get information about a patient on demand while supporting a distributed model of storage.

Each participating entity will be required, at a minimum, to provide two interfaces, as shown in Figure 2. The first is a Web service will receive a query including a patient identifier or locator as an input from the outside world, and will return the health records stored at that facility to the requesting entity. This Web service will pool all of the data located within that entity to construct the return record in a specified format.

Additionally, each entity in the network will need a Web service application that consumes the other participants' search requests. This service returns the information to the requester, which then decides what to do with the information received. However, if an additional entity is added to the exchange, including the Web services described above, the original participating entities will need to update their applications to consume the new Web services provided by the new entity. Using this approach eventually leads to every entity having to

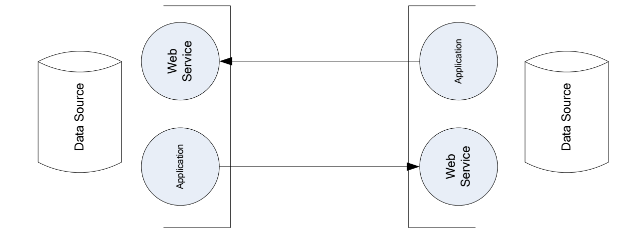

Figure 2. Basic information exchange between two entities. Each entity provides a Web service that allows the other entity to access information.

have a separate application to consume everyone else's Web services.

Instead of updating all applications to consume a new Web service each time that a new entity joins in the exchanging of information, as shown in Figure 3, we suggest using a neutral Web service that acts as a registry of all entities wishing to exchange data. To participate in the information exchange, an entity registers its Web services with the neutral Web service. Then, instead of directly consuming Web services of each individual entity participating in the exchange, a participating entity consumes the neutral Web service, sending requests and information through it. A search request sends a patient locator to the neutral Web service, which then searches all of the registered entities' Web services for the relevant data and returns the data as an XML Document within a SOAP envelope to the requester. Note that the neutral service does not store the information, but rather acts something like a router to move requests and responses to the appropriate entity's interface.

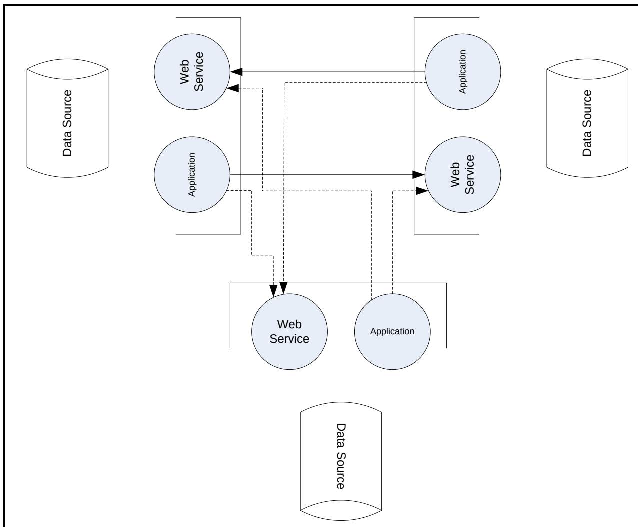

Figure 3 Adding a new entity to an exchange system. The solid lines are established data exchanges. The dotted lines are data exchanges that would need to be established when a new entity joined the data exchange.

In order to provide Web services that will receive and answer these messages, each entity will need a server capable of hosting Web services. This Web server should be able to cross the firewalls between the local EHR and the Internet in a secure, controlled manner. This server could provide connections to the local EHR either through database drivers that enable the Web service to query the local EHR system directly, or an interface that supports the consumption of Web services internally that can return the requested information.

Using this approach, we propose that an interoperable health information network supporting the exchange of information among separate entities, rather than through a centralized database, can be implemented as shown in Figure 4. A standardized information interface for each facility in the network is created using Web services. When a facility needs access to information that is stored at another location, its local system issues a request that is sent to a central management system aware of all entities participating in the network. The management system then forwards the request to the Web servers at the appropriate location(s) and routes the response back to the request's originator.

This design is flexible, in that the originating location need not itself be aware of all other participating entities. The central management system acts like a router in determining where the request should be sent. If the system includes some kind of registry that can identify which locations have relevant information, then the requests would only be sent to those locations. Alternatively, the request could be sent to all locations in the network, with the understanding that a null/negative response will be returned by some of them.

This system also protects the independence and privacy of each individual location, as the Web services interface determines what information from its internal records is included in the response to a request. Not all information in a record need be released, so locally sensitive information remains within the local record and only the information explicitly approved for release is sent back to the requestor.

Obviously, for this approach to work, a standardized set of messages will need to be defined. This could be based on HL7 or another approach. Once these standards have been determined, each participating entity in the information exchange network will be able to create a set of Web services that will accept and return messages. Note that the exchange network does not necessarily need to define the format of these messages: rather, it must standardize conventions about what information will be exchanged and what data types will be included. It is important to note that, as long as a data type is not proprietary to a particular system, data of that type can be exchanged among different languages and operating systems easily.

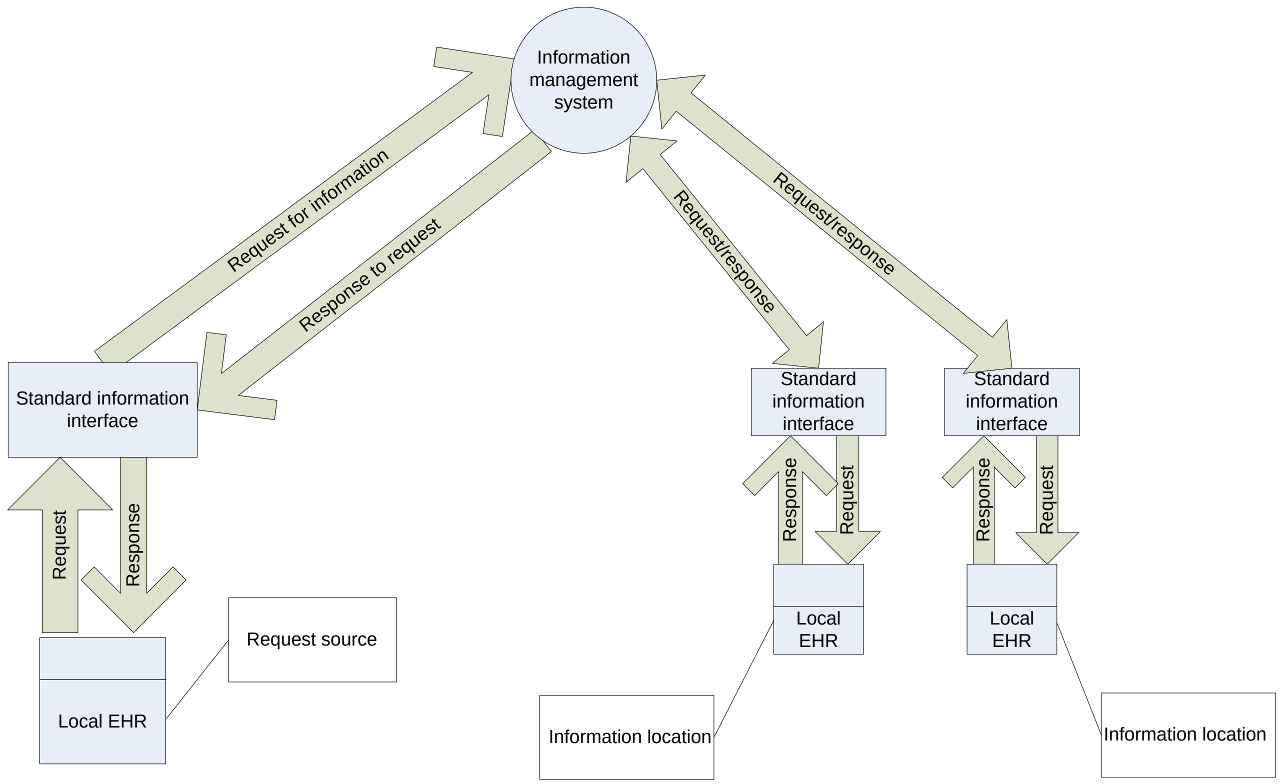

Figure 4. A system for interoperable exchange of health information.

4.0 Conclusion

Health care institutions have been investing in computerized systems for many years. What is different about the current environment is that the health care industry is no longer just considering automation of the administrative or back-office work within an institution. These new information systems are intended to allow communication across systems and institutions for a broader array of purposes. By collecting electronic health data from multiple sources and analyzing or presenting them in usable form at the point of care or in other contexts, an interoperable health information system has the potential to improve the quality and safety of care delivery, enhance the development and dissemination of evidence-based protocols, or allow more effective allocation of health care resources.

At the same time, the development of health information systems poses several challenges. Mandates for interoperable health information exchange do not address the issues of disparate systems, incompatible data formats, and changing standards. For these reasons health care providers are in a holding pattern and are either waiting for further standardization before addressing the issue, or not even entering the EHR arena until these issues have been solved.

However, technical developments over the last decade have made it possible for disparate systems within a facility to exchange information. XML and SOAP have provided the means to exchange information among these systems. The introduction of Web services takes advantage of these technologies to integrate these systems with a mechanism for interoperability.

Using Web services to exchange information among health care providers is a natural step forward. Web services deal well with changing standards by allowing each provider to make changes to one set of services, rather than to several different interfaces. With standards of information exchange in a state of flux, this approach allows adding or changing information within the Web services relatively easy.

In addition, adding a new participant to the collaboration has a minimum impact to the whole system in this approach. At a minimum, only the neutral Web service will need to be changed to consume the new Web service. It may even be possible to create a dynamic linkage to new services, so that only a registration function is needed to maintain the neutral service.

References

[1] R. N. Charette. Dying for Data: A comprehensive system of electronic medical records promises to save

lives and cut health care costs - but how do you build one? *IEEE Spectrum* 43(10):22-27, 2006.

[2] National Electronic Health Records Taskforce 2000, *A Health Information Network for Australia*, Commonwealth Department of Health and Aged Care, Canberra.

[3] *HL7 Organization: Health Level 7*. Retrieved 11/13/2006. http://www.hl7.org.

[4] Ean-Wen Huang, Da-Wei Wang, Der-Ming Liou. Development of a deterministic XML schema by resolving structure ambiguity of HL7 Messages. *Computer Methods and Programs in Biomedicine* 80:1-15, 2005.

[5] K. Sung Um, Y. Sik Kwak, H. Cho, and I Kon Kim. Development of an HL7 interface engine, based on tree structure and streaming algorithm, for large-size messages which include image data. *Computer Methods and Programs in Biomedicine* 80(2):126-140, 2005.

[6] M.A. Cupito. Systems integration: Revving up the interface engines. *Health Management Technology* 18(13):18-20, 1997.

[7] H. Deitel, P. Deitel, B. DuWaldt, L.K. Trees. *Web Services: A Technical Introduction.* Upper Saddle River, NJ; Prentice Hall, 2003.