Virtual Consolidation of Boston's Beth Israel and New England Deaconess Hospitals via the World Wide Web

John D. Halamka MD and Charles Safran MD, MS
Center for Clinical Computing, Beth Israel Deaconess Medical Center, Boston, MA

With the advent of Integrated Healthcare Delivery Systems, medical records are increasingly distributed across multiple institutions. Timely access to these medical records is a critical need for healthcare providers. The CareWeb™ project provides an architecture for World Wide Web-based retrieval of electronic medical records from heterogeneous data sources. Using Health Level 7 (HL7), web technologies and readily available software components, we consolidated the electronic records of Boston's Beth Israel and Deaconess Hospitals. We report on the creation of CareWeb™ (freya.bidmc.harvard.edu/careweb.htm) and propose it as a means to electronically link Integrated Health Care Delivery Systems and geographically distant information resources.

INTRODUCTION

In an era of increasing competition for health care dollars, medical institutions are merging and consolidating with increasing frequency. Given that the majority of such institutions have heterogeneous hospital-based computing resources, integrating information systems across merged institutions is a difficult problem.

The CareGroup was formed in 1996 by the merger of the Beth Israel Hospital, the Deaconess Hospital, three Boston area community hospitals, and several satellite outpatient clinics, creating a billion-dollar integrated healthcare delivery system. A major post-merger issue has been the integration of existing electronic medical records. Each site has different legacy systems, different institutional vocabularies and varying completeness of clinical information.

The CareWeb™ project was conceived to provide a means for the virtual consolidation of the medical records at these heterogeneous institutions.

CareWeb™ is an actual implementation of the W3EMRS architecture [1]. The W3EMRS architecture provides the first comprehensive approach to delivering a unified virtual medical record for both inpatient and outpatient care settings. With W3EMRS, a request for information is made via a web browser to an "agglutinator". The agglutinator requests information from several heterogeneous hospital, clinic and office sites. A site server is implemented for each site which interprets the agglutinator request and returns the requested information as an HL7 data stream. The agglutinator merges data from all responding sites and returns a unified web page to the web browser. The W3EMRS project is an example of a World Wide Web-based approach that has delivered promising results in a laboratory environment.

On-going research continues to support the suitability of web technology for clinical applications, although deployed systems in a clinical environment are rare. Several authors have explored the use of the web for display of medical records for a single institution [7-8] . Others have proposed multi-institutional approaches utilizing a central repository [5-6], updated at fixed intervals, which limits the timeliness of data.

CareWeb™ is the first project to use the World Wide Web to successfully consolidate heterogeneous clinical data across multiple institutions. CareWeb™ implements web-exposed HL7-based [2] medical information servers at each participating institution in the healthcare delivery network. A central "Consolidator" processes requests for information from healthcare providers and queries all sites on the network. The Consolidator then delivers an integrated multi-institutional medical record to the health care provider.

Creating such a system presented many issues, both technical and political, including definition of a common medical record, creation of a master patient index strategy and implementation of a security/confidentiality policy.

METHODS

Design Considerations

We extended the W3EMRS architecture in several ways. For security, we implemented hardware tokens and RSA cryptography. For the software

1091—8280/97/$5.00

implementation, we created an object-oriented architecture. The necessary software tasks were organized into modular functionality and standard, commercially available software objects were obtained to provide this functionality. Objects included an HL7 messaging product, an HTTP communication product, and legacy system interfacing products. By choosing available software objects, we both minimized our development time and leveraged the expertise of others who encapsulated their specific knowledge into the generation of these components.

Defining the Common Medical Record

To create an integrated medical record across heterogeneous institutions, we first needed to define the types of data we wanted to share.

Some institutions, such as the Beth Israel hospital, have complete online medical records with detailed notes for most visits and procedures [11]. Others, such as the community hospitals, currently keep only basic demographic and visit information on-line.

The power of the W3EMRS/CareWeb™ concept is its ability to provide a consolidated view of the existing records of multiple institutions. To maximize the information displayed, we designed a common medical record that includes patient demographics, medications, allergies, visits and notes.

Strategies for the Master Patient Index

Much work has been done on the details of creating master patient indexes [3]. Examining the types of demographic data available across the CareGroup institutions, we elected to use two master patient index strategies.

For competent patients who are able to communicate identifying information, we elected to use patient last name, patient first name, date of birth, gender, mother's first name and father's first name to identify patients uniquely. Although not perfect, we believe that this strategy will suffice for more than 90% of our patients.

For incompetent or critical patients, we elected to use patient last name, patient first name, date of birth and gender. Although potentially not precise enough to identify patients uniquely, we believe that this strategy will maximize the information available to a health care provider about an otherwise non-communicating patient.

ARCHITECTURE

The clinical data at the Beth Israel Hospital is stored in a comprehensive, custom built MUMPS based system composed of 28,000 programs. The clinical data at the Deaconess Hospital is stored in a Sybase clinical data repository.

CareWeb™ site servers operate behind the web servers of each hospital and create a link to the underlying legacy systems at each institution. These site servers interpret incoming HL7 requests for information, translate them into specific legacy system queries and package the resulting information into an HL7 response.

To allow users to query multiple hospitals simultaneously, we developed a CareWeb™ "Consolidator", which processes user requests, dispatches them to multiple hospitals' site servers, and processes the information retrieved (Figure 1).

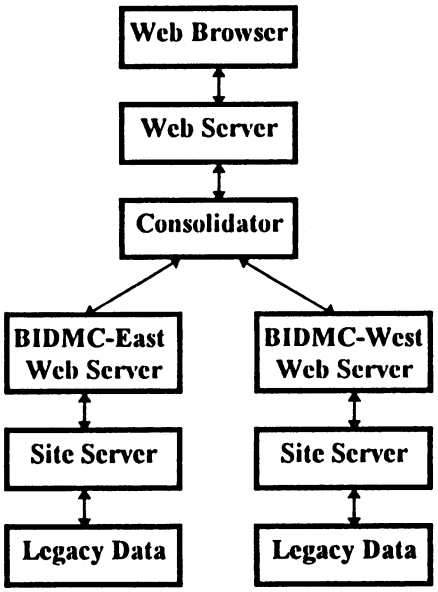

Figure 1 The CareWeb™ Architecture

A typical session begins when a health care provider on a standard web browser creates a query for information by specifying patient identification. This information is submitted via standard HTML forms to the Consolidator. The Consolidator generates an HL7 request for information to both the Beth Israel and Deaconess site servers. The site servers return HL7 encoded demographics, problems, medications, allergies, notes and visits. The Consolidator interprets the incoming messages and creates a single unified presentation which it

sends back to the health care provider as a series of web pages. Full navigational control is enabled with tool bars that allow the medical record to be scanned using a tab folder-like paradigm.

The CareWeb™ project was implemented using Visual Basic complemented by ActiveX components for HL7 messaging, HTTP communications, and legacy system interfacing. The Consolidator and site servers were created as multithreaded, interrupt-driven Internet Information Server Active Server Pages [4].

Site Servers

We defined HL7 semantic mappings for the CareWeb™ common medical record (patient demographics, medications, allergies, visits and notes) and created site servers which populated these mappings via interfaces to each institution's legacy systems.

For the Beth Israel Hospital, we explored several strategies for connecting site servers to the Center for Clinical Computing MUMPS data structures. Using the Intersystems VisM tools for Visual Basic, we built a robust interface between Visual Basic and MUMPS legacy data. Visual Basic provided a flexible interface to Internet Information Server and VisM enabled our legacy system queries.

Using KB Systems KB_SQL, we implemented Open Database Connectivity (ODBC) access to MUMPS data. This ODBC link enabled not only Visual Basic but other ODBC compliant tools such as Microsoft's Active Data Objects Components to directly query MUMPS data structures.

For the Deaconess, we used the Microsoft SQL Server Open Database Connectivity (ODBC) Driver to connect to the Deaconess' Sybase Clinical Data Repository.

Site servers were implemented in Visual Basic as OLE Automation Servers and their methods were called by Microsoft Internet Information Server's Active Server Pages.

Consolidator

The Consolidator interprets the incoming health care provider request, translates the request into an HL7 query, sends the query via the HTTP post method to each site server, receives the site server response, and consolidates the collected results into a single presentation.

The Consolidator was written in Visual Basic as an OLE Automation Server called by Microsoft Internet Information Server's Active Server Pages. Consolidator HL7 messaging was performed by an HL7 ActiveX component. HTTP messaging was likewise handled by an HTTP ActiveX component.

To implement the presentation layer we elected to use simple HTML 3.0 without ActiveX or Java browser side components. This decision was made to insure browser independence. Given the diversity of machines distributed throughout the CareGroup, we wanted CareWeb™ to perform equally well on Mosaic as with the latest versions of Internet Explorer or Netscape.

SECURITY AND CONFIDENTIALITY

Security was a primary concern in the creation of CareWeb™. Web-enabling medical records raises many legitimate concerns for potential breaches of patient confidentiality. Working with members of the W3EMRS collaborative and with Beth Israel/Deaconess staff, we created an initial security architecture.

User Authentication

We guarantee the authenticity of a user by using hardware tokens. A user must be in possession of a hardware device in order to access pages containing patient data. Hardware tokens are a robust way to provide not only security but also user authentication. Most tokens are small, handheld devices containing microprocessors that calculate and display unpredictable codes. These codes change at a specified interval, typically 60 seconds. Utilizing Security Dynamics SecurID tokens, Microsoft's Active Server Pages and VBScript, we implemented a method for manipulating encrypted browser-side cookies to maintain authentication and authorization information.

Securing browser to server communications

To create security between a browser running on a user's desktop and the Consolidator web server, we implemented the Netscape standard Secure Sockets Layer [9]. The SecurID username and passcode are only exchanged after an encrypted connection has been established by the Secure Sockets Layer.

Securing consolidator to site communications

For communications between the Consolidator and site servers, we implemented RSA public key encryption for key exchange, session key

cryptography for data exchange and digital signature for confirmation of the veracity of the Consolidator request [10]. Each Consolidator HL7 request is signed with the Consolidator's RSA private key. The request is sent to the site server which uses the Consolidator's public key to validate the digital signature through standard hashing and signature verification methods. The site server retrieves the information requested and signs the HL7 response with its private key. The site server then generates a session key which it uses to encrypt the HL7 response. The session key is then encrypted using the Consolidator's public key. The encrypted session key and encrypted data are sent back to the Consolidator. The session key is decrypted using the Consolidator's private key. The encrypted HL7 response is decrypted using the decrypted session key. Finally, the HL7 response is validated using the site server's public key. All decrypted site server messages are consolidated into a single web page and returned to the original requesting browser over the Secure Sockets Layer (figure 2).

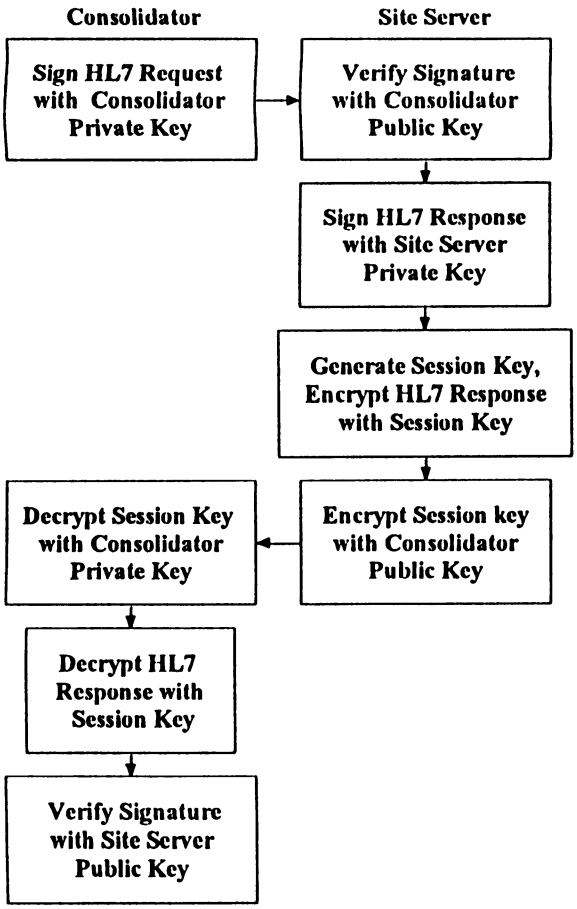

Figure 2 Encryption of Network Transmissions

Securing Legacy System Data

The existing legacy systems at the Beth Israel and Deaconess hospitals employ a complex series of hardware controls which limit internet transactions from outside the institution. Using routers and firewalls, network administrators limit legacy system access to hardware devices physically located within the campuses.

Auditing

It has been the security policy of the Beth Israel hospital to provide auditing at the level of the specific patient queried and the individual menu selections used. CareWeb™ provides auditing at the site server level to track specific requests made of each institution, including date, time, type of request, requester, and institutional patient id queried. These audit trails may be used to discover inappropriate access patterns and to provide patients with a record of accesses of their medical record, should patients request such information.

EVALUATION

As an early evaluation of the CareWeb™ architecture, we sought and received Institutional Review Board (IRB) approval to web-expose selected medical records from actual patients who have records at more than one CareGroup institution. Patient approval was obtained and patients were allowed to view the CareWeb™ versions of their medical record before making them generally available. Furthermore actual patient names and addresses were pseudonymized, but medical information was not altered.

The system was evaluated by 25 healthcare providers, chosen at random from both institutions, who assessed CareWeb's ease of use, response times and utility in patient care. Further evaluation was performed by 25 information systems staff members who evaluated CareWeb's robustness, security and potential for deployment in the live environment.

During the evaluation period, the CareWeb™ system processed 3000 accesses for patient information. 23/25 healthcare providers rated the utility of patient information for emergency care as excellent and 2/25 good. 25/25 ranked ease of use as excellent and 25/25 judged response times as fast. 24/25 information systems staff members judged its live deployment potential to be excellent, 1/25 good. 23/25 judged robustness to be excellent, 2/25 good.

25/25 judged security and confidentiality measures to be excellent.

Further evaluation in a live environment is planned over the next six months. CareGroup is currently extending internet services to all of the Emergency Departments in its healthcare delivery network. After institutional approval, CareWeb™ will be made available as an Emergency Department resource.

DISCUSSION

During the development of CareWeb™, we encountered many technical challenges which we overcame by adhering to existing standards and maintaining an object oriented architecture.

Although each institution in the CareGroup has varying completeness of clinical information, we found that the CareWeb™ architecture was able to handle inconsistencies among institutions by displaying common medical record information from each institution in site-native vocabularies, on a single web page. Additional sites may be added to the Consolidator without program modification, resulting in easy extensibility.

Our early evaluation of CareWeb™ demonstrates that its implementation is useful, robust, and provides rapid response times.

Several political challenges still await resolution before the architecture is widely adopted to provide time critical data to a regional or nationwide network of care providers. Continuing reports of flaws in internet security give a public impression that the web is not a suitable environment for sensitive information. We believe we have addressed security and confidentiality issues with the CareWeb™ security model.

Extending CareWeb™ beyond a single integrated delivery system will result in data sharing by competing networks of providers. Web-exposing clinical data for access by competitors will likely be met with resistance by senior management. During our continued deployment of the CareWeb™ system, we will have the opportunity to address the ongoing political issues created by live implementation of the system. Further plans for the CareWeb™ system are focused on extending the variety of health data objects available. CareWeb™ architectures which include laboratory results, medical images and multimedia objects are currently under development.

ACKNOWLEDGEMENTS

We gratefully acknowledge the mentorship of Isaac Kohane, Peter Szolovits, David Rind, and the other members of the W3EMRS working group.

Funded in part by a cooperative agreement with the Agency for Health Care Policy and Research and the National Library of Medicine Sharing Paperless Records among Networks of Providers (U01- 08749), and the Douglas Porter Fellowship, Center for Clinical Computing.

1. Kohane, I, van Wingerde FJ, Fackler JC, et. al Sharing Medical Records Across Multiple Heterogeneous and Competing Institutions, Proceedings 1996 Fall AMIA 608-612.
2. Health Level Seven: An application protocol for electronic data exchange in healthcare environments, version 2.2, Chicago, Illinois Health Level Seven 1990.
3. Duncan K, Patient Indexing Strategies, ComNets, February/March 1996.
4. Microsoft, Active Server Pages White Paper, Microsoft, 1996.
5. National Information Infrastructure Health Information Network, Essential Emergency Dataset, NIIHIN Consortium,Mequon,WI ,1995.
6. Pollock D, Data Elements for Emergency DeparTMent Systems Draft 3 Dataset, Center for Disease Control and Prevention, 1996.
7. McDonald CJ, Overhage JM, Tierney WM et al, The RMRS: Physician use for input and output and web browser based computing, Proceedings 1996 Fall AMIA, 989.
8. Willard KE, Hallgren JH, Sielaff B, Connelly DP, The Deployment of a World Wide web (W3) Based Medical Information System, Proceedings 1995 Fall AMIA 771- 775.
9. Hickman KEB, Elgamal T, The Secure Sockets Layer Protocol 3.0, Internet Draft, Netscape Communications Corporation 1996
10. Schneier H , Applied Cryptography, John Wiley and Sons, 1996
11. Safran, C, Role of Computing in Patient Care in Two Hospitals, MD Computing 6:3 1989.