MEDINFO 98
B. Cesnik et al. (Eds)
Amsterdam: IOS Press

CareWeb™, a Web-based Medical Record for an Integrated Healthcare Delivery System

John D. Halamka MD and Charles Safran MD, MS

Center for Clinical Computing, Harvard Medical School and
the Department of Medicine, Beth Israel Deaconess Medical Center, Boston MA

Abstract

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.

Keywords
Patient Records; Internet; Security

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 implementation of the W3EMRS[1] architecture, which uses the World Wide Web to 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.

Materials and Methods

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).

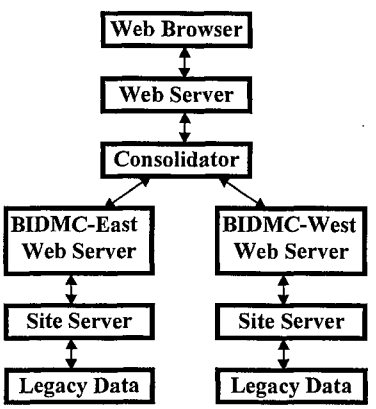

Figure 1 - The CareWeb Architecture

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.

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 [3].

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

In March of 1997, the National Research Council (NRC) of the National Academy of Sciences issued the report "For the Record: Protecting Electronic Health Information"[4] Concluding that the current practices at the majority of health care facilities in the United States are insufficient, the Council delineated both technical and organizational approaches to protecting electronic health information. We incorporated all thirteen NRC recommendations into the CareWeb™ security architecture [8,9].

Strong Enterprise-wide Authentication

We guarantee the authenticity of each user with Security Dynamics SecurID hardware tokens. These tokens are small, handheld devices containing microprocessors that calculate and display unpredictable codes. These codes change at a specified interval, typically 60 seconds. Our implementation requires that each user accessing CareWeb™ begin a session by entering a username, a memorized personal identification number (PIN) and the currently displayed password from the SecurID device. This information is transmitted to a security server which authenticates the user and verifies that the correct password was entered. The security server compares the user-entered password with its knowledge of what password should have been entered for that 60 second period. If the password does not match, it also checks the password from the previous 60 second period to account for delays in typing and transmission. Once a password is verified, the user is authenticated for the entire enterprise for the duration of the web session or 15 minutes, whichever is less. An encrypted security "cookie" is sent back to the user's browser and this cookie is automatically used for all future security dialogs. Using Visual Basic Script and Microsoft's Active Server Pages, we dynamically decrypt the cookie within the web server and invisibly re-verify authentication before responding to additional requests for healthcare data.

If the SecurID token is lost or stolen, it can be immediately deactivated for the entire enterprise by disabling it at the security server.

Access Validation

In addition to storing encrypted username and password information, the security cookie contains the job role of the user. Displays of healthcare information are generated dynamically by Active Server page scripts, which assemble the multi-institutional medical record. The scripts can tailor delivered health care information based on the job role indicated by the cookie. In our proof-of-concept implementation, we have restricted this tailoring of access to specific areas of the medical record such as discharge summaries. We have not created a facility to scan for and restrict specific content *within* an area, such as removing a psychiatric evaluation from a discharge summary.

Expanded Multi-organizational Audit Trails

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 [5]. CareWeb™ implements a complete *multi-organizational* audit trail.

In any multi-institutional architecture there are two places to capture the audit - either at the institutional level where the information is stored (the sites) or at the point where the information is delivered (the CareWeb™ "Consolidator"). We elected to capture the information at the site level. Although

only a single CareWeb™ "Consolidator" exists today, CareWeb™ could be expanded such that other regional or national "Consolidators" might query information from the CareGroup institutions. If the audit was captured at the "Consolidator" level, each institution would have to rely on the security practices of the "Consolidator" operators. By storing audit trails at each site, each hospital can control and audit the information that leaves its site, regardless of where it is delivered. Each hospital site server captures patient identification information, requester, the requester's IP address, date, time, and information requested.

Although information is stored at the site level we have implemented a multi-institutional auditing system that provides patients with the details of the movement of their medical information throughout the healthcare enterprise. The auditing query system has the same hardware token authentication and access controls as are required for any CareWeb™ healthcare data request. Once authenticated, an auditor enters patient identification information and submits the information to an "Auditing Consolidator". This "Auditing Consolidator" uses secure, password protected Open Database Connectivity (ODBC) connections to query the audit trails of the individual hospitals. It produces a consolidated report showing all flows of information about the patient for all institutions.

Protection of External Communications

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 campus.

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 [6]. The SecurID username and passcode are only exchanged after an encrypted connection has been established by the Secure Sockets Layer.

Encryption of public network transmissions

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 authentication of the Consolidator and site servers [7]. 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.

Electronic authentication of records

The use of hardware tokens for system access also facilitates electronic signature. Since possession of the hardware device authenticates the user, the SecurID token is used as the official electronic signature for "signing" all CareWeb™ documents and audit trails.

As noted above, digital signature cryptography methods are used for all network transmissions, ensuring the integrity of all health data delivered. The NRC recommends an implementation of hashing and digital signature to insure that medical records are not changed on the individual systems where they are stored. In the CareWeb™ architecture we have no control of the integrity of the data stored at each institution. We have created a secure mechanism to transport each institution's data and can guarantee that the data was not changed during the retrieval process. The reputability of the data is dictated by security policies of each institution providing the data.

Physical security and Disaster Recovery

The notion of a multi-institutional architecture provides significant physical protection for health data. Instead of physically locating all patient records in a central data source which is vulnerable to physical disasters, the CareWeb™ architecture depends upon the consolidator which stores no health care information. All that is needed to restore a physically destroyed "Consolidator" system is to connect another computer containing the "Consolidator" software and its required cryptographic keys to the hospital network. Currently, all site servers are geographically dispersed and are locked in secure computer rooms accessed by electronic keycode. In the CareWeb™ architecture we have no control of the physical security and disaster recovery practices of the individual sites which provide data. However, if any sites sustain a disaster and cease to provide data, the Consolidator notes that a site is currently unavailable and provides a virtual medical record comprised of all functioning sites.

Software discipline

No browser software is installed on either the site servers or the Consolidator machines, precluding inappropriate downloads. Virus checking programs are in place on all CareWeb™ systems and are executed daily by a system daemon.

On the end-user workstation, we have been careful not to cache pages returned by the Consolidator. In our laboratory environment we have verified that neither Netscape nor Internet Explorer cache pages that have been returned via a secure socket connection such as that used by CareWeb™. We cannot protect against an authenticated user who installs a new type of browser that does cache secure pages. However, all pages returned by the Consolidator have an HTML header, which indicates that they expire on delivery. Even if a new browser was installed which cached information, this expiration forces the browser to replace each cached page as new requests for information are made, minimizing the amount of information that is stored on the end-user workstation.

System Assessment

Daily assessment is performed on both the Consolidator and site server systems. On the Consolidator, a security log lists all SecurID tokens used, all failed login attempts, and all changes made to the token database. Web server log analysis (WebTrends) shows all attempts to contact the Consolidator web server showing IP address, time, date and page accessed. System assessments are also performed on a daily basis at each institutional site, per their own institutional guidelines.

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 heathcare 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.

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 collaboration of Isaac Kohane, Peter Szolovits and David Rind.

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 P. Porter Fellowship, Center for Clinical Computing.

References

[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] Microsoft, Active Server Pages White Paper, Microsoft, 1996.

[4] For the Record: Protecting Electronic Health Information, Computer Science and Telecommunications Board, National Research Council, National Academy Press, 1997

[5] Safran C, Rind D, et al, "Protection of Confidentiality in the Computer-based Patient Record", MD Computing Vol 12, No. 3, 1995.

[6] Hickman KEB, Elgamal T, The Secure Sockets Layer Protocol 3.0, Internet Draft, Netscape Communications Corporation 1996

[7] Schnieier H , Applied Cryptography, John Wiley and Sons, 1996

[8] Rind D, Kohane I, Szolovits P, Safran S, Chueh H, Barnett G, Maintaining the Confidentiality of Medical Records Shared over the Internet and World Wide Web, Annals of Internal Medicine, July 1997 (in press).

[9] Halamka J, Rind D, Safran C, "A WWW Implementation of National Recommendations for Protecting Electronic Health Information", JAMIA, November 1997 (in press).