Application of Information Technology ■

A WWW Implementation of National Recommendations for Protecting Electronic Health Information

JOHN D. HALAMKA, MD, PETER SZOLOVITS, PHD, DAVID RIND, MD, CHARLES SAFRAN, MD, MS

Abstract 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." 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. The Beth Israel Deaconess Medical Center recently implemented a proof-of-concept, Web-based, cross-institutional medical record, CareWeb, which incorporates the NRC security and confidentiality recommendations. We report on our WWW implementation of the NRC recommendations and an initial evaluation of the balance between ease of use and confidentiality.

■ **J Am Med Inform Assoc.** 1997;4:458–464.

In his 1997 State of the Union address, President Clinton noted that "we should connect every hospital to the Internet, so that doctors can instantly share data about their patients with the best specialists in the field."[1] The security and confidentiality implications of Web-connecting the nation's clinical data are a major impediment to realizing this noble goal.

In 1995, the National Research Council (NRC) of the National Academy of Sciences was charged with evaluating the practical measures that can be used to reduce the risk of improper disclosure of confidential health information while providing justified access to those interested in improving the quality and reducing the cost of care. Their March 1997 report, "For the Record: Protecting Electronic Health Information," presents the findings of 2 years of collaborative investigations and site visits.[2]

The report reviews the public policy context, as well as the internal and external threats to organizations that possess health information, and it outlines technical and organizational approaches to protecting health information.

We implemented all of the technical recommendations of the report in the context of CareWeb, a proof-of-concept, Web-based, multi-institutional medical record that integrates the Beth Israel and New England Deaconess hospitals. Creating such a system presented many challenges, both technical and political. The implementation is displayed on the Web at the following address:

http://freya.bidmc.harvard.edu/careweb.htm

Background

The NRC recommendations are separated into two categories: eight technical practices for immediate implementation and five practices for future implementation. This discrimination is made with the assumption that the immediate practices will suffice as a minimum for the current state of technology.

Affiliations of the authors: Department of Medicine, Beth Israel Deaconess Medical Center and the Center for Clinical Computing, Harvard Medical School, Boston, MA (JDH, DR, CS); Laboratory for Computer Science, MIT, Cambridge, MA (PS).

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

Correspondence and reprints: John D. Halamka MD, Center for Clinical Computing, 350 Longwood Avenue, Boston, MA 02115. E-mail: jhalamka@bidmc.harvard.edu

Received for publication: 5/22/97; accepted for publication: 7/14/97.

However, as more health information is available in electronic form and as more security technology becomes generally available, the need for more complete security implementations becomes necessary.

Practices for Immediate Implementation

Individual Authentication of Users. The NRC site visits discovered that many health care organizations have generalized login usernames/passwords such as MD for physicians and RN for nurses. To properly authenticate individuals on any computer system containing health care data, every individual should have a unique username/password for access. Such a policy allows individuals to be held accountable for all actions taken while logged on.

Access Controls. Many health care computing systems allow all users to view all information. There is, however, no good reason for a laboratory technician to read the confidential full text data contained in a patient psychiatric profile. Health care providers should be allowed to view clinical information on a need-to-know basis. The most obvious implementation of such controls would be to assign access to different health care computing functions based on job role.

Audit Trails. Although newspaper articles highlight the threat of computer break-ins by unauthorized "hackers" from outside health care organizations, inappropriate health care data access from inside the organization is far more common. Normal human curiosity leads individuals not involved in a patient's care to look up the records of VIPs, celebrities, and fellow employees. If authenticated users are to be held accountable for actions taken while using the health care computing system, retrievable audit trails that log all accesses to information should be kept. These logs should include time, date, information accessed, and user ID. Audit trails should be available for patient review on demand.

Physical Security and Disaster Recovery. Computer terminals should be positioned where they cannot be accessed by unauthorized users. Unauthorized personnel must be denied access to paper printouts and electronic storage. Backup tapes should be made frequently, and tapes should be housed off site in the case of a physical disaster.

Protection of Remote Access Points. Firewalls should be implemented to provide strong, centralized security and to allow outside access to only those systems critical to outside users. All remote accesses should be protected by single session or encrypted passwords.

Protection of External Electronic Communications. All patient-identifiable data transmitted over public networks should be encrypted.

Software Discipline. Virus-checking programs should be installed on all servers, and downloads from the Internet to servers should be limited.

System Assessment. Audits should be performed on a monthly basis to examine vulnerability to password cracking programs and to verify procedures implemented to detect system vulnerabilities.

Practices for Future Implementation

Strong Authentication. Health care providers occasionally share usernames/passwords or write them down near a computer terminal. Such practices defeat the authentication, access controls, and audit trails offered by unique usernames/passwords. Authentication is substantially strengthened by requiring that logon be paired with physical possession of "hardware tokens," such as smartcards, magnetic strip IDs, or devices with constantly changing passwords.

Enterprise-wide Authentication. Health care environments typically have many heterogeneous computing systems. If users are forced to have different logon information for each computer system, remembering such information is an inconvenience, and users will tend to write down login information. To minimize such behavior, users should authenticate once and then have access to all relevant systems.

Access Validation. In the simplest form of access control, different system functions are available based on job role. A more sophisticated implementation would tailor content within functions by job role. For example, a discharge summary could be viewed by both a physician and a billing coder, but details of the patient's psychiatric evaluation would not appear for the coder.

Expanded Audit Trails. Simple audit trails capture information at a single organization. An expanded auditing system would provide interorganizational audit trails that trace information as it passes thorough the health care complex.

Electronic Authentication of Records. An electronic signature should be used to "sign" submitted medical records, and a cryptographic digital signature should be used when retrieving records to ensure records are not modified during the transmission process.

System Description

We implemented all 13 NRC recommendations in the CareWeb system.

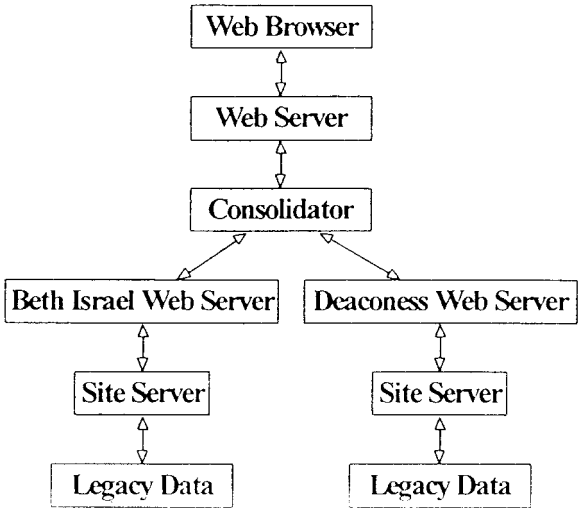

Figure 1 The CareWeb architecture.

The CareWeb 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 unites these systems using an implementation of the W3EMRS architecture.[4,9]

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 HL-7[5] requests for information, translate them into specific legacy system queries, and package the resulting information into an HL-7 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 (Fig. 1).

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 HL-7 request for information to both the Beth Israel and Deaconess site servers. The site servers return HL-7 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.

Security Architecture

To implement the NRC recommendations, we reduced the present and future recommendations to nine action areas: strong enterprise-wide authentication, access validation, expanded audit trails, protection of external communications, encryption of public network transmissions, electronic authentication of records, physical security and disaster recovery, software discipline, and system assessment.

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 health care 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 health care information are generated dynamically by Active Server page scripts, which assemble the multiinstitutional 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.[6] CareWeb implements a complete multiorganizational audit trail.

In any multiinstitutional 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 so that other regional or national Consolidators might query information from the CareGroup institutions. If the audit were captured at the Consolidator level, each institution would have to rely on the security practices of the Consolidator operators. If audit trails are stored at each site, each hospital can control and audit the information that leaves its site, regardless of where the information 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 multiinstitutional auditing system that provides patients with the details of the movement of their medical information throughout the health care enterprise. The auditing query system has the same hardware token authentication and access controls as are required for any CareWeb health care 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 that 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.[7] 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.[8] This is shown in Figure 2.

Each Consolidator HL-7 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 HL-7 response with its private key. The site server then generates a session key, which it uses to encrypt the HL-7 response. The session key is then encrypted, using the Consolidator's public key. The encrypted

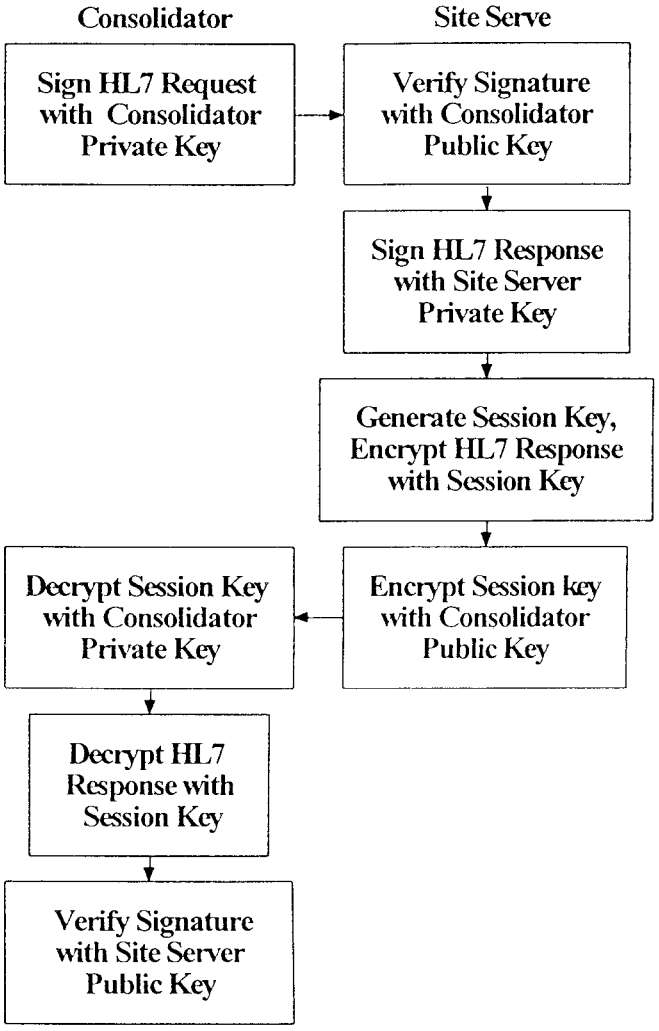

Figure 2 Encryption of public network transmissions.

session key and encrypted data are sent back to the Consolidator. The session key is decrypted using the Consolidator's private key. The encrypted HL-7 response is decrypted using the decrypted session key. Finally, the HL-7 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 implementing hashing and digital signatures to ensure that medical records are not changed on the individual systems where they are stored. In the CareWeb architecture, we have no control over 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 integrity of the data is dictated by security policies of each institution providing the data.

Physical Security and Disaster Recovery. The notion of a multiinstitutional architecture provides significant physical protection for health data. Instead of physically locating all patient records in a central data source 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 that 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 composed of all functioning sites.

When deployed in the production environment, the personal computers used to perform CareWeb lookups will be located at emergency department nursing stations. Screens will be specifically turned away from care areas, as is the standard practice of the medical center.

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 that indicates that they expire on delivery. Even if a new browser were installed that cached information, this expiration would force the browser to replace each cached page as new requests for information were 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, displaying IP address, time, date, and page accessed. System assessments are also performed on a daily basis at each institutional site, according to their own institutional guidelines.

Status Report

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 records before making them generally available. Furthermore, pseudonyms were used for actual patient names and addresses, but medical information was not altered.

The security architecture was implemented in one man-month using standard Microsoft Windows NT architectures and ActiveX components costing under $10,000.

The proof-of-concept security architecture implementation was Web-exposed, and the Web site processed 3,000 requests for health care information. Fast re-

sponse times and high reliability were evidenced by the fact that 100% of transactions were completed on the first request. Reviewing our audit trails and security logs, we found neither unauthorized access nor inappropriate denial of access to the site.

The system was evaluated by 25 health care 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. Initial reactions to the prototype appear to be positive.

Further evaluation in a live environment is planned over the next 30 days. CareGroup is currently extending Internet services to all of the emergency departments in its health care delivery network. CareWeb is being deployed as an emergency department resource.

Discussion

The political sensitivity of Web-exposing confidential data was emphasized in March 1997, when the Social Security Administration created a publicly accessible Web page for display of social security benefits information. The page was discontinued in April 1997 because of an outcry from privacy advocates and citizens' rights groups.[3]

To protect our Web-based medical record system, we implemented all 13 of the National Research Council's recommendations, including those suggested for future implementation. Our architecture includes several innovations for protection of health care data. These include the use of hardware tokens for authentication, the use of cryptographic methods for protection of information flows over public networks, and the creation of a multiorganizational auditing system.

Although individual elements of the proof-of-concept CareWeb security architecture have been implemented elsewhere, CareWeb provides a unique multiinstitutional approach. Building on the Beth Israel tradition of providing patient-focused services, we have created a multiinstitutional auditing architecture that gives the patient a comprehensive view of all lookups performed in an integrated health care delivery system. Such a view maximizes the meaningfulness of the audit trail from the patient's viewpoint, as it provides a snapshot of patient information flow throughout the health care enterprise.

We are currently mobilizing the resources to transition CareWeb from a proof-of-concept system to a production system. Issues include administration of security hardware tokens; providing 24 hours per day, 7 days per week support; and obtaining all necessary institutional approvals.

CareWeb provides a rapidly deployable, low-cost architecture that leverages the strengths of existing institutions to form an integrated health care delivery system. Current limitations of the architecture are that it does not address the validity of data provided by the individual hospital sites, it relies on the physical security and disaster policies of each site to protect site-specific data, and it makes no attempt to duplicate patient data at a central location to use as a secondary resource in case of site failure.

Health care organizations face many threats, both internal and external, to the confidentiality of health care information.[6,9–11] Technical and organizational practices are needed to ensure that critical health information is always available to authorized users and denied to those without the need to know. The balance between ease of use and confidentiality requires careful implementation. Denial of appropriate access to patient records imposed by a failure of the security technology can lead to disastrous health care consequences.

Although convenience concerns were raised by a minority of health care providers surveyed, the majority of health care providers interviewed were satisfied with the CareWeb implementation.

To date, the evaluation of the CareWeb security architecture has been limited to the proof-of-concept system. An expanded evaluation of the deployed version will include detailed feedback from over 50 daily users of the system.

Any technical implementation must be complemented by a strong organizational policy to sanction those who inappropriately access health care data. The Beth Israel Deaconess Medical Center has a long-standing tradition of protecting patient confidentiality within its legacy systems, and organizational policies are already in place. This greatly facilitated the implementation of the CareWeb security architecture and minimized the barriers to implementation that would be found at institutions without such existing policies.

The Health Insurance and Portability and Accountability Act of 1996 (Kennedy–Kassebaum) requires that the Secretary of Health and Human Services submit to Congress, by August of 1997, detailed recommendations on protecting the privacy of individually identifiable health information. The CareWeb security architecture offers an early trial implementation of several potential strategies.

Conclusion

Our experience with the CareWeb system has demonstrated the feasibility of using the Web to allow access to longitudinal patient record data distributed across multiple sites, providers, and institutions. We have demonstrated that a security architecture can be built around this system to provide a balance between allowing ease of access to emergency health care data and protecting patient confidentiality. This security architecture builds on the work of others[4,9] to create the first Web-based implementation of the National Research Council's recommendations for present and future security practices.

We have learned many valuable lessons during the development of CareWeb. First, the organizational barriers to deploying a secure Web-based medical record can outweigh the technical challenges. Continuing reports of flaws in Internet security give a public impression that the Web is not a suitable environment for sensitive information, and this creates difficulty in obtaining institutional support. Consensus for deploying such a system must include information systems personnel, hospital administrators, public relations specialists, and the clinicians themselves.

Second, the existing hospital infrastructure provides a strict limitation on the types of technology that may be deployed. Although the Beth Israel Deaconess Medical Center has a campus-wide, high-speed network infrastructure, many machines are incapable of running the current versions of Web browsers, preventing the use of Java and browser-side scripting languages.

CareWeb is currently being deployed in the live production environment, and we will report on the challenges encountered.

References ■

1. Clinton, William. State of the Union Address, February 4, 1997.
2. For the Record: Protecting Electronic Health Information, Computer Science and Telecommunications Board, Commission on Physical Sciences, Mathematics and Applications, National Research Council. National Academy Press, 1997.
3. Schwartz J, Saffir BJ. Privacy concerns short-circuit Social Security's online service. Washington Post, April 10, 1997.
4. Kohane I, van Wingerde FJ, Fackler JC, et al. Sharing medical records across multiple heterogeneous and competing institutions. Proc. AMIA Annu Fall Symp. Philadelphia: Hanley & Belfus, 1996;608–12.
5. Health Level Seven: An application protocol for electronic data exchange in healthcare environments, version 2.2. Chicago: Illinois Health Level Seven, 1990.
6. Safran C, Rind D, et al. Protection of confidentiality in the computer-based patient record. MD Comput. 1995;12:187–92.
7. Hickman KEB, Elgamal T. The Secure Sockets Layer Protocol 3.0, Internet Draft. Netscape Communications Corporation, 1996.
8. Schnieier H. Applied Cryptography. New York: John Wiley & Sons, 1996.
9. 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. Ann Intern Med. 1997;127:138–141.
10. Woodward B. The computer-based patient record and confidentiality, N Engl J Med. 1995;333:1419–22.
11. Barrows RC, Clayton PD. Privacy, confidentiality and electronic medical records. J Am Med Inform Assoc. 1996;3:139–48.