The Practice of Informatics

Review ■

Privacy, Confidentiality, and Electronic Medical Records

RANDOLPH C. BARROWS, JR., MD, PAUL D. CLAYTON, PHD

Abstract The enhanced availability of health information in an electronic format is strategic for industry-wide efforts to improve the quality and reduce the cost of health care, yet it brings a concomitant concern of greater risk for loss of privacy among health care participants. The authors review the conflicting goals of accessibility and security for electronic medical records and discuss nontechnical and technical aspects that constitute a reasonable security solution. It is argued that with guiding policy and current technology, an electronic medical record may offer better security than a traditional paper record.

■ **JAMIA.** 1996;3:139–148.

One purpose of electronic medical records (EMRs) is to increase the accessibility and sharing of health records among authorized individuals. Privacy of information collected during health care processes is necessary because of significant economic, psychologic, and social harm that can come to individuals when personal health information is disclosed.[11,34,38] With remote access to distributed health data, or the pooling of health data from multiple sites in a central repository, the potential for loss of information privacy is greater than in isolated EMR systems, or in systems with paper medical records, when proper safeguards are not taken. With appropriate safeguards, however, computer-based medical records may actually offer more security than traditional paper-record systems. Applicable security technologies exist and have proved effective in the banking and military sectors, but experience is lacking to ascertain whether current technologies are satisfactory for health care. As yet, no model security implementations exist in any clinical computing environment,[34] although awareness of risks and of possible technical solutions is increasing.

In this review, we examine the extent to which fears of the loss of privacy due to EMRs are justified, and we discuss measures to protect the security of health data. We also consider the trade-offs between accessibility and security of EMRs compared with paper records.

Goals of Informational Security in Health Care

A cohesive informational security policy is lacking across institutions, counties, and states, and governmental and nongovernmental committees are grappling with difficult policy details that have far-reaching consequences. Although the establishment and implementation of security policies may be challenging, the goals of information security in health care can be simply stated[10,11,20]:

1. To ensure the privacy of patients and the confidentiality of health care data (prevention of unauthorized disclosure of information)

Affiliation of the authors: Department of Medical Informatics, Columbia University, New York, NY.

Correspondence and reprints: Randolph C. Barrows, Jr., MD, Center for Medical Informatics, Columbia Presbyterian Medical Center, 1310 Atchley Pavilion, 161 Fort Washington Avenue, New York, NY 10032. e-mail: barrows@cucis.cis.columbia.edu

Received for publication: 11/2/95; accepted for publication: 11/8/95.

2. To ensure the integrity of health care data (prevention of unauthorized modification of information)
3. To ensure the availability of health data for authorized persons (prevention of unauthorized or unintended withholding of information or resources)

The goal of information privacy raises issues of access control (user authentication and authorization) and the application of cryptographic protocols for data transmission and storage. The goal of data integrity introduces the need for electronic user and data authentication.[9,21] The goal of data availability raises issues of access control, system reliability, and backup mechanisms (system and data redundancy). The policy and technical aspects of these and related issues are discussed below.

Security Policy

As many others have pointed out,[1,3,11,34,39] the main problem with information security in health care is not technology, but a lack of cohesive security policy. Policy must shape technology, not vice versa. Security policy defines what is to be protected, to what reasonable degree protections will be afforded, and who is privileged to access protected items. A policy is influenced by:

1. The functional requirements of an information system (what users need to accomplish from the system)
2. The security requirements for the system (items that need to be protected)
3. A threat model (the expected motives and resources of potential perpetrators)

The role of policy is to balance the functional and security requirements of a system, which are typically at odds. Security requirements can often be tempered by the practical concerns of a threat model, because costs and user inconveniences rise sharply with harsher security implementations.

"Inside attacks,"[17] the most routine kinds of security transgressions, represent one example of a threat concern. Such attacks are committed by persons who are legitimate system users with privileges but who abuse their privileges in search of gossip material, or for other personal or financial motivations. The monetary value of health data obtainable on most individuals, however, is relatively low (unlike some financial data or military secrets), so it is reasonably safe to assume that an attacker will not spend inordinate resources (money and time) on attempting to acquire such data by computer break-in or cryptanalytic attack. Specifically desired information, as always, might be available with less trouble and expense via "social engineering" techniques (bribery, extortion, personal misrepresentation of identity, and so forth). The health data of celebrities and other prominent persons may be of greater monetary value in certain markets, but currently available (although not necessarily implemented) security mechanisms, such as system management, access control, and encryption techniques, are sufficient to thwart or detect the covert activities of hospital employees, newspaper reporters, relatives, and other unsophisticated attackers.

Another example of potential threats comes from information-hungry employers, insurance companies, and managed care organizations. These organizations have greater economic resources, along with the motivation of significant profit from what they can know about individuals. Unethical operations in such industries could allocate a high-end computer to the task of breaking a cryptographic key used in the transmission of health data over inexpensive public channels. The 1995 cost of a machine capable of breaking a Data Encryption Standard of the U.S. government (DES) key within 1 year (with an 8% chance per month) is only $64,000.[35] Profit-motivated health care–related organizations and unethical "private investigators" might be willing to make this investment and, for example, gather HIV data, which could be used on a covert basis to deny medical insurance coverage.

The above threats concern attacks on patient privacy, but threat models should also consider attacks on the integrity and availability of health data. Such threats might come from malevolent "hackers," natural disasters, or mechanical failures and could potentially cost data guardians more than any breach in confidentiality.

The Data Security Policy and Standards developed for the Mayo Clinic/Foundation provide one model example of a clear institutional security policy statement.[27] As an example of an approach to policy setting, Columbia-Presbyterian Medical Center (CPMC) hired external consultants to facilitate security policy development for its Integrated Advanced Information Management System project.[6] After 24 meetings with 80 people from numerous departments that spanned two institutions, 14 overlapping topic areas for which policy development was needed were identified:

1. User authentication—issues relating to the iden-

tification of a user to the system and the ways in which the system might know that a user is who they claim to be.

2. Physical security of data center sites—issues relating to the physical access to computer hardware; theft prevention; backup and disaster recovery; and the security of sensitive terminal locations, such as console or control, and of publicly accessible terminals.
3. Access control to system resources—issues of the physical devices and logical mechanisms, such as computer programs, that control access to system resources.
4. Data ownership—issues of who will own which data, the delegation of authority over data, and enunciation of the duties and responsibilities of data ownership.
5. Data protection policies—issues of minimally acceptable and consistent protections to be afforded by systems crossing organizational and functional boundaries, anticipated implementation barriers to those protections, and the punitive measures for organizational members abusing system privileges.
6. Building security into systems—issues of how to assure that security requirements are addressed in central and local participating systems, how to partition security responsibilities between central and local systems, and how to assure that security requirements remain satisfied as systems are modified or expanded.
7. Security of hard copy materials—issues of how to prevent security breaches from paper copies of sensitive electronic documents and data.
8. Systems integrity—issues related to the accuracy and reliability of system data, and the integrity and reliability of physical computer and network systems.
9. User profiles—issues related to defining user types and roles that serve to distinguish the functional needs and security levels of users.
10. Legal and liability issues—issues relating to the uses and misuses of the system that involve potential liabilities or legal concerns for participating organizations, including protections under existing computer crime laws, liabilities when a record is compromised, and requirements for user penalties under union contracts.
11. Problem identification and resolution—issues of system audits and auditability, intrusion detection and notification of intrusions, and detection and notification mechanisms for other types of security problems.
12. Network security—issues relating to the security management of computer networks and the movement of data over such networks, including the security of bridges and routing equipment, the passing of authorization tokens, data encryption, electronic signatures, and nonrepudiation of messages.
13. Informed consent—issues related to the use of medical information collected about patients and obtaining consent from patients for desired and potential uses of medical data.
14. Education of users—issues related to the education of users regarding their responsibilities as system users and the risks conjured by their actions, including activities on the system and degrees of nonvigilance.

From these 14 areas, a list of 65 policy items needing definition were identified. These items were then ranked, resulting in a list of 17 urgent actions. Of particular note, the number one action item was to establish a mechanism for making institutional policy.

Privacy and Confidentiality in Health Care

The relationship between health care provider and patient is one characterized by intimacy and trust, and confidentiality is embedded at least implicitly in patient-provider interactions. The notion of confidentiality in health care has a strong professional tradition that has suffered progressive erosion due to third-party reimbursement schemes, managed care and other health care organizational structures, and the perceptions and culture of professionals within modern health care systems.[24] One third of medical professionals have indicated that information is given to unauthorized people "somewhat often."[38]

Unfortunately, information privacy has an incomplete and inconsistent legal basis.[15,28] Federal law prohibiting information disclosure pertains only to information associated with federal agencies, not to information held in the private sector or by state and local governments. Most states have laws that address at least minimally the privacy of medical records but do not consistently recognize computerized records as legitimate documents.

One reason for the difficulty in setting policy is that the legal concept of privacy is relative and shifts from time to time to reflect the public versus private interests of society.[33] Consider, for example, current airline-passenger and baggage-inspection policies compared with those of 30 years ago, and laws that require the reporting of infectious, especially sexually transmitted, diseases. In addition, privacy is partly in the eye of the beholder, and an intrusion of privacy perceived by one person may be considered as a convenience by others (targeted marketing, mail-order catalogs, solicitations by insurers and service-providers of preventive health, and so forth).

In a 1993 survey, 80% of persons believed that consumers had lost control over information about themselves.[38] EMR developers should strive to maintain the confidentiality of personal health information to foster public trust in information systems that hold promise for improving health care quality and decreasing the costs of care. For their own benefit and the benefit of society, patients should not be made reticent in sharing medically relevant information with health care practitioners.

The goal of strict information privacy conflicts with goals of optimal patient care, however, as well as with medical research, public health, and social policy, all of which may require access to patients' confidential medical records without their explicit knowledge or consent. In addition, health care providers have a working need for high data availability and are intolerant of cumbersome security procedures. For instance, when access hurdles are too steep, logon sessions and passwords may be shared among providers. Because the use of information technology in health care is still relatively new and not yet ubiquitous, there is generally too little awareness of the risks conjured by such actions.

Technically, the confidentiality medical records in computers can be maintained proactively by both access-control mechanisms and audit trail logs (discussed below), which can be inspected proactively or in response to suspicious events. Other mechanisms for assuring confidentiality include the education of EMR users regarding security concerns, professional responsibilities, and personal accountability; time-outs on system terminals; hard-copy control; clear policies; and consistent disciplinary actions. Human factors, however, such as errors, negligence, and unethical activities, can result in breaches of confidentiality despite optimal security implementations.

Accordingly, the American Civil Liberties Union (ACLU) believes that a privacy policy for health information should be based on the following principles[18]:

1. Strict limits on access and disclosure must apply to all personally identifiable health data, regardless of the form in which the information is maintained.
2. All personally identifiable health records must be under an individual's control. No personal information may be disclosed without an individual's uncoerced, informed consent.
3. Health-record information systems must be required to build in security measures to protect personal information against both unauthorized access and misuse by authorized users.
4. Employers must be denied access to personally identifiable health information on their employees and prospective employees.
5. Patients must be given notice of all uses of their health information.
6. Individuals must have a right of access to their own medical and financial records, including rights to copy and correct any and all information contained in those records.
7. Both a private right of action and a governmental enforcement mechanism must be established to prevent or remedy wrongful disclosures or other misuse of information.
8. A federal oversight system must be established to ensure compliance with privacy laws and regulations.

Pending federal legislation with bipartisan support (the "Bennett Bill")[41] seeks to implement recommendations to protect the confidentiality of medical information and to guarantee access to patients of their own health data, with the hope that such measures will promote a health-information infrastructure. The bill has drawn sharp criticism, however, from consumer-rights advocacy groups like the ACLU due to lack of patient controls over how personal health information may be used and disseminated, particularly regarding the compilation of health information within certified "health information services."[37]

The Joint Commission on Accreditation of Health Organizations has begun to demand that patients' rights, security policies, and information-management standards be addressed in more explicit ways.[31] The 1995 standards proposed significant new requirements in these areas. In recognition that most health care organizations are not yet able to meet those standards, the 1996 version downsized the information management chapter by more than 70 requirements,[42] with the stated intention of a more gradual deployment.

Data Ownership and Legal Accountability

Data ownership is a legally complex issue. Ownership of a medical record is at best a limited right that is primarily custodial in nature, and information contained in the record is often characterized as the patient's property.[16] Any immediate and clear legal assignment of electronic health data ownership, from which may follow assignment of responsibility, does not appear likely. All parties who are entrusted with health data, both the movers and the users, should reasonably be considered as stewards of that data, and may be held liable for irresponsible acts and breaches of confidentiality.

Informed Consent to Disclosure

An informed consent to disclosure of information typically requires that the patient:

1. Be told what information is to be disclosed.
2. Understand what is being disclosed.
3. Is competent to provide consent.
4. Consents willingly, free from coercion.

Implementation of the doctrine of informed consent to disclosure involves many potential difficulties, and "informed consent," as it pertains to the typical uses of health care data, is arguably a misnomer. Infirm or confused patients cannot meaningfully sign an informed release, and no informed release specifically covers all potential or desired uses of medical data that may be collected on an individual. Also, patients are coerced into giving up personal rights to confidentiality when they apply for insurance or sign a hospital waiver that allows medical information to be shared. In recognition of such concerns, a general release of medical information in New York state no longer applies to HIV data. Finally, patients are typically asked to authorize disclosure of medical information, yet only about half of the states guarantee a patient's right to see his or her own medical record.

Traditionally, patients have difficulty gaining access to their own records, and without knowledge of what is contained in the record, consent for disclosure cannot be fully informed. The position of the American Health Information Management Association reflects a balance of opinion and states that an EMR requires that patients have greater access to their own medical record.[5] The proposed Bennett Bill would guarantee that right, except when disclosure might endanger the life or safety of any individual, or information in the record identifies a confidential source of information about the requesting patient.[41]

Use of Medical Data

The established primary uses of medical records are in providing health care, paying for it, and assuring its proper delivery. Secondary uses of medical data include those made by various business and governmental organizations such as life and auto insurers, employers, licensing agencies, public health agencies, the media, medical researchers, educational institutions, rehabilitation and social welfare programs, and uses for legal purposes. Responsibility for the protection of patient privacy and the confidentiality of computerized medical information must extend to these secondary users. Institutional policy should dictate how patient data may be used and to whom information will be released.

When electronic records are used for research, valid epidemiologic studies may be conducted using aggregates of nonidentifiable patient data. The Bennett Bill requires specific patient authorization when such "scrubbed" data are inadequate.[41] In addition, encrypted patient identifiers might provide acceptable research results and still adequately protect patient privacy.

User Authentication and Access Control

Originators of the few landmark computer-based patient-record systems have grappled with the aforementioned conflicting goals of security and functionality in health care systems.[7,40] Usually, systems use some form of password security for user authentication, and user-specific or role-specific menus may be used to implement further limitations on access. However, standard password access controls do not prevent insider threats and are not helpful when authentication has been compromised.

In addition, tight access control at the level of the type of user, computer application, or patient fails in critical ways in the health care environment.[9,10] Sensitive data (i.e., mental health data or HIV status) are often among the most important items necessary to take care of a patient. This is the information that may need to be made available and shared among numerous care providers and ancillary health personnel. Most often, numerous persons at multiple levels in multiple roles (medical students, residents, nurses, therapists, dietitians, social workers, administrators, consultant physicians, covering physicians, and a private or personal "attending" physician) are routinely involved in a patient's care, and it is difficult to predict which person in which role will validly need access to a person's health record at some particular time. Provisions for emergencies, when none of the patient's usual care team is around, must also be made. Thus, in an EMR setting, prohibition of access by most medical users to most data on most patients is often not practical. For this reason, clinical system pioneers have usually allowed all clinical personnel access to the computerized medical record of all patients in a hospital, and often to the records of patients not in the hospital as well (i.e., records of discharged patients or their ambulatory care, or both).

Improved multilevel and role-based access models for health care that better accommodate user needs are under development.[8,12,22,23] A "need-to-show" model (versus the military "need-to-know" multilevel security model) and its supportive technical platform have been proposed, with the specific intention of extending the notion of individual professional accountability for health data to interaction with information systems.[29] Such accountability may help discourage information sharing across unauthorized informal human networks,[13] a problem that is difficult to address by technology.

The determination of how much effort should go toward authenticating a person is a matter of institutional policy. User identifiers with password authentication are often employed, but other technical solutions, such as biometric authentication by morphometric hand measurements or voiceprints, system-synchronized random-number generating cards, and passphrase-encrypting smartcards, are more expensive, but they may be more effective alternatives when deemed compatible with policy considerations.

As an example of an approach to access control, the CPMC Clinical Information System (CIS) implements an access-control matrix with one axis representing user roles (attending physicians, residents, medical students, hospital nurses, clinic nurses, various types of technicians, and so forth) and the other axis representing data types (laboratory data, radiology reports, discharge summaries, demographic information, and so forth). We defined 68 user types and six classes of data. Departmental leaders make the determination of access privileges for each user type, subject to the approval of the hospital medical board. Users receive a menu of options specific for their defined access privileges. Login screens remind users that information is limited to legitimate medical purposes and that misuse can lead to dismissal as well as civil and criminal penalties. Access to data on VIPs and hospital employees invokes an additional screen message warning that all user activities are recorded. A similar approach at Boston's Beth Israel Hospital, along with a system utility that allows users to review the names of persons who have looked at their electronic record, was reported to effectively deter "insider" abuse of system privileges.[40]

Cryptography

Cryptographic techniques applicable to the goals of privacy, integrity, and access control have not yet been significantly deployed in the health care environment, and experience is needed before establishing that they could provide security solutions compatible with the diversity of health care needs.[19]

As a trivial example of an encryption cipher, the famous Caesar Cipher uses a "shift-by-three" rule, so that every "A" in a message is replaced by a "D," every "B" by an "E," and so forth. The algorithm is said to have been used by Julius Caesar to encode communications with his generals via human messengers whom he did not trust. Many more complicated and secure mathematical algorithms for encryption exist. Private-key, or "secret-key," encryption depends on a number or string of characters that is shared only between the communicating parties and is used by an encryption algorithm to encode and decode the message. The exact encryption algorithm need not be a secret. The best known such encryption algorithm is DES, mentioned above. A main problem with private-key encryption protocols is that communicating parties must somehow securely share and use the "secret" key.

The use of public-key encryption can avoid some of the pitfalls of the need to share a secret key by making use of a mathematical technique that creates an "asymmetrical cryptosystem," that is, the keys to encode and decode a message are different but intimately linked, so that they are, in effect, functional inverses of each other and can only be used together. In public-key cryptography, one key is published, and the other remains private to a user. To send a secret message, the sender obtains the recipient's public key and uses it to scramble the message, which the recipient can decode with his or her private key. In addition, the creator of a message or document can "sign" it by encoding a piece or algorithmic "digest" of the document with his or her secret key, so that anyone can then verify the "signature" by decoding it with the signer's published key.

The New York State Community Health Management Information System (NYSCHMIS) Confidentiality and Data Security Policy says:

> All data collected into or handled through the repository and defined as 'deniable' (identifiable) ... shall be encrypted, both when being transmitted through the network or if written to a local system. Software and/or hardware shall be supplied with secure algorithms which will encrypt/decrypt all such sensitive data.[32]

For practical purposes, due to the imbedding of sensitive data in text documents, we recommend that all health data in an EMR environment be encrypted when transmitted over public or insecure channels and when residing on storage devices in local machines.

The Massachusetts Institute of Technology's Kerberos is a secret-key cryptographic protocol for the provision of authentication and authorization services in a distributed environment. Although its use has been outlined for the health care setting, it has not been implemented.[9] Public-key cryptographic protocols have been proposed to address the need for a patient identifier that is universal (across institutions and states).[36] Software tool kits for the secure transmission and archiving of files by medical applications are beginning to appear.[21] In the near future, vendor products will supply encryption technology embedded within computer systems for health care. Until then, EMR developers are forced to create their own implementations of well-known and secure cryptographic algorithms and protocols.[35]

Data Integrity

Electronic patient data can be assumed valid based on software testing and verification, access-control mechanisms, and error-checking protocols used in data transport, or they can be additionally authenticated as valid with digital signatures, as discussed above. Most lapses in data integrity will continue to be due to human error and to malfunctions or "bugs" in medical computer systems.

Firewalls

Firewalls are computers that are positioned between a site's internal network and an unsecured public network, such as the Internet, and may be useful at EMR sites. Firewall computers are configured to monitor and regulate the messages passing into and out of a site's private network and so can prevent unauthorized users from entering local computer systems from the outside, or can prevent particular programs and services from operating through the firewall. Such functionality can help protect private information from leaving an EMR site, or can impose an extra layer of password security on authorized users.

Reliability, Redundancy, and System Backups

As discussed above, threat models should consider potential "attacks," whether accidental or intentional, on the integrity and availability of health data. Hardware or software failures, including "denial-of-service" attacks, can cause downtime or loss of vital health care data for EMR users. The reliability of EMR systems and data should be considered a security concern and should be covered in security policy and system management activities, usually through mechanisms that support data redundancy and system backups.

Audit Trails

Primarily because of limitations on the applicability of access-control methods in health care, the audit trail has become a critical tool for managing issues of data security. In any large computing environment is a significant risk for misuse of the system by authorized users. For this reason, the audit trail has become an important reactive security mechanism and is often used for post hoc detection of security violations and for support of disciplinary actions.

For example, at CPMC, the CIS records both the identity of any individual who looks at patient data and the type of data accessed. In one illustrative instance, a resident physician (physician in specialty training) in obstetrics harassed a nurse about being pregnant before the nurse had announced her pregnancy to any individual. The nurse complained, and review of audit-trail data showed that the resident physician had indeed looked at the nurse's test results, and without a valid "need to know," this led to an official reprimand.

One problem with audit-trail data is that the data are typically far too voluminous for human processing. "Level C2" is a U.S. Department of Defense computer security classification requiring auditing and the unavailability of encrypted passwords, and a level C2 audit mechanism for a multiuser system can fill 1 gigabyte of disk space within an hour.[30] One published prototype system generated 7 megabytes (MB) per day per average user, and up to 136 MB per busy user.[2] The CIS audit-trail logs as implemented at CPMC fill about 100 MB of disk space per month. Typically, 95% of audit data are of no security significance,[4] and use of the data accumulated in security audit files is at best minimal. Extraneous data in the files obviously makes it harder to detect suspicious behavior, especially that which might be detected by complex relationships between the data features, something particularly difficult for humans to discover.

Automated reduction and analysis tools for audit trail data could help immensely, but their availability has been limited. Frank discusses data-reduction methods for intrusion detection and gives an example of selection methods used to identify a subset of data features that best classify some audit data.[25] Systems that implement some kind of automated analysis of audit-trail data are a relatively recent development. Early approaches to audit-trail analysis only categorized threats as due to internal versus external penetrators, but the current goal is to identify threats by any users or processes that attempt an illegal action within their authorized boundaries (abuse of system privileges), or that attempt an action not within their authorized boundaries (exceed system privileges), as well as any action by unauthorized system users, such as intruders that masquerade as authorized users or otherwise evade system authentication and security controls.[26] Later models for performing intrusion detection have used statistical user profiling or expert system techniques that examine the deviation of actual user behaviors from anticipated or usual behaviors on the system.[14]

One way to distinguish intrusion-detection methods is based on the type of intrusion: anomaly detection versus misuse detection.[30] Misuse detection involves well-defined patterns of intrusion that exploit weaknesses in software and can be detected directly. Because it searches for known vulnerabilities, misuse detection is of little use in detecting new or unknown intrusive behaviors. Anomaly detection depends on unusual behavior or unusual use of system resources, and it seeks to detect the complement of normal behavior. In general, intrusive activity is expected to be some subset of anomalous activity; however, intrusive behavior does not always coincide with anomalous behavior and might be accomplished as the sum of individual nonanomalous activities.

Nine developed intrusion detection tools are reviewed by Marshall.[4] Most of these systems perform both anomaly and misuse detection. Statistical techniques lend themselves to anomaly detection but are inadequate to detect all types of intrusions and do not prevent users from gradually training their usage profiles, so that activity previously considered anomalous might be regarded as normal. Expert systems and model-based techniques lend themselves to misuse detection, but specification of the orderings on facts, for the pattern matching of events, has been deleteriously inefficient.[30] Thus, in the best systems, anomaly and misuse detection methods complement each other.

Each system is out of necessity however, somewhat ad hoc and custom designed. Few systems are general or flexible enough to be easily portable or adaptable. More generic systems, capable of reuse and retargeting, are likely to be inefficient or of limited power. Also, the cost of building an intrusion-detection system is high and requires specialized knowledge input from system and security experts who can make an appropriate choice of statistical metrics and can specify expert rules. Moreover, testing and validation of intrusion-detection systems are difficult, because potential attack scenarios can be difficult to simulate, and the lack of a common audit-trail format precludes easy comparisons between the performance of existing systems and common attack scenarios.

Consequently, no commercially available audit-analysis tool kit exists, and there is as yet no known application of software tools for audit analysis in the health care sector. The idea, however, was discussed by Shea and colleagues[9] and is apparently under active implementation in the European community.[26]

A Comparison of the Paper and Electronic Record Environments

Many security issues discussed to this point can apply to paper-based as well as electronic records. The most obvious new risk factor afforded by the electronic records is also the benefit that pushes us toward the electronic format: enhanced convenience of accessibility and distribution of health information. A related and potentially troubling capability is the ability to query for a population of patients who have a common feature (such as, the same surgeon or a particular test result). Any risks of an electronic breach of security must be weighed against analogous risks and recognized disadvantages of paper record systems. Electronic records are arguably more secure if the proper policies and best available technologies are in place.

For example, paper medical records do not allow one to obtain an accurate audit trail of who has seen the record and what portions of the record were accessed. Also, the use of paper records make it difficult to restrict certain classes of users to see only particular types of information. Paper records are easily altered by removal or substitution of documents, but an electronic document signed with an encrypted digital signature is much more difficult to alter. The paper record can be in only one place at a time, whereas the same information in electronic format can be available to multiple users simultaneously. Also, the content of the computer-based medical record can be presented in a clearly organized and legible fashion, so that caregivers will more likely respond to important information. In a paper-based environment, real-time rule-based suggestions and warnings cannot be generated when standards of health care are missed. Also, costs may soon favor the use of electronic record systems. For example, at CPMC, the cost to find and pull a paper record from the file room for doctors, for just a single patient visit, has been estimated to be between $5 and $10. In contrast, we estimate that the total cost for the creation and lifetime maintenance of an electronic record for our patients is between $25 and $50.

Thus, substantial advantages to the electronic record exist, and it seems prudent to move ahead with implementations of electronic records, including the policies required to guide the application of available security technologies.

Conclusion

Although security concerns surrounding health data in EMR environments are justified, solutions are surmountable with currently available technologies. In the banking industry, analogous security implementations have allowed greater personal convenience, including access to personal bank accounts from a choice of locations and at all times of day, without security compromises. Although neither automatic bank tellers nor electronic medical records are free from instances of abuse, implementation of available protocols for electronic systems probably provides better security than the security measures that are used in analogous manual systems. In any security system, the weak links are most likely to be human.

A major challenge will be that of enticing developers, who are eager for working medical computer applications, to make the financial and time investments in designing and building adequate security features into their systems. Institutional policies will be a key stimulus in this regard. Chief financial officers will likely come to regard security investments as insurance policies: although we must pay for the policies, we are pleased when there is no need to file a claim. A more formidable barrier than security requirements to the implementation of sharable records in an EMR environment is the current lack of convenient and acceptable ways to acquire data from patients and providers in an electronic format. Security issues should not deter progress toward solving this more substantial problem.

References ■

(Sorted Chronologically by Year and
Alphabetically by First Author
within Each Year)

1968

1. Curran WJ, Stearns B, Kaplan H. Privacy, confidentiality and other legal considerations in the establishment of a centralized health-data system. N Engl J Med. 1968;281:241–8.

1987

2. Picciotto J. The design of an effective auditing subsystem. Proceedings of the 1987 IEEE Symposium on Security and Privacy. Washington, DC: IEEE Computer Society Press.

1990

3. Brannigan V, Beier B. Standards for privacy in medical information systems: a technico-legel revolution. In: Miller RA, ed. Proceedings of the Fourteenth Annual Symposium for Computer Applications in Medical Care. Los Alamitas, CA: IEEE Computer Society Press, 1990:266–70.

1991

4. Marshall VH. Intrusion detection in computers: a summary of the Trusted Information Systems (TIS) Report on Intrusion Detection Systems (TIS report #348). McLean, VA: Booz, Allen & Hamilton, January 29, 1991.

1992

5. American Health Information Management Association. Position Statement. Chicago: AHIMA, March 1992:1.
6. Clayton PD, Sideli RV, Sengupta S. Open architecture and integrated information at Columbia-Presbyterian Medical Center. MD Comput. 1992;9:297–303.
7. Murphy G. System and data protection. In: Ball MJ, Collen MF, eds. Aspects of the Computer-based Patient Record. New York: Springer-Verlag, 1992:201–13.
8. Orr GA, Brantley BA. Development of a model of information security requirements for enterprise-wide medical information systems. In: Frisse ME, ed. Proceedings of the Sixteenth Annual Symposium for Computer Applications in Medical Care. New York: McGraw-Hill, 1992:287–91.
9. Shea S, Sengupta S, Crosswell A, Clayton PD. Network information security in a Phase III Integrated Academic Information Management System (IAIMS). In: Frisse ME, ed. Proceedings of the Sixteenth Annual Symposium for Computer Applications in Medical Care. New York: McGraw-Hill, 1992:283–6.

1993

10. Bakker AR. Security in medical information systems. In: van Bemmel JH, McCray AT, eds. Yearbook of Medical Informatics. New York: Shattauer, 1993:52–60.
11. Gostin LO, Turek-Brezina J, Powers M, Kozloff R, Faden R, Steinauer DD. Privacy and security of personal information in a new health care system. JAMA. 1993;270:2487–93.
12. Henkind SJ, Orlowski JM, Skarulis PC. Application of a mulilevel access model in the development of a security infrastructure for a clinical information system. In: Safran C, ed. Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care. New York: McGraw-Hill, 1993: 64–8.
13. Lincoln TL. Privacy: a real-world problem with fuzzy boundaries. Methods Inf Med. 1993;32:104–7.

14. Lunt LT. A survey of intrusion detection techniques. Comput Security, 1993;12:405–18.
15. U.S. Government, Office of Technology Assessment. Medical Privacy Report, 1993. Chapter 1: Introduction, Summary and Options. Washington, DC, 1993.
16. U.S. Government, Office of Technology Assessment. Medical Privacy Report, 1993. Chapter 3: Computerized Health Care Information. Washington, DC, 1993.
17. U.S. Government, Office of Technology Assessment. Medical Privacy Report, 1993. Appendix A: Selected Topics in Computer Security. Washington, DC, 1993.

1994

18. American Civil Liberties Union. Toward a New Health Care System: The Civil Liberties Issues. An ACLU Public Policy Report (ISBN 0-914031-24-4); New York, February 1994.
19. Barber B, Bakker A, Bengtsson S. Conclusions and recommendations. Int J Biomed Comput. 1994;35(suppl 1):221–9.
20. Bengtsson S. Clinical requirements for the security of the electronic patient record. Int J Biomed Comput. 1994;35(suppl 1): 29–31.
21. Bleumer G. Security for decentralized health information systems. Int J Biomed Comput. 1994;35(suppl 1):140–5.
22. Brannigan VM. A framework for "neéd to know" authorizations in medical computer systems: responding to the constitutional requirements. In: Ozbolt JG, ed. Proceedings of the Eighteenth Annual Symposium on Computer Applications in Medical Care. JAMIA. 1994 suppl:392–6.
23. Dargahi R, Classen DW, Bobroff RB, et al. The development of a data security model for the collaborative social and medical services system. In: Ozbolt JG, ed. Proceedings of the Eighteenth Annual Symposium on Computer Applications in Medical Care. JAMIA. 1994 suppl:349–53.
24. France FH, Gaunt PN. The need for security—a clinical view. Int J Biomed Comput. 1994;35(suppl 1):189–94.
25. Frank J. Artificial Intelligence and Intrusion Detection: Current and Future Directions. Davis, CA: Division of Computer Science, University of California-Davis, June 9, 1994.
26. Hayam A. Security audit center—a suggested model for effective audit strategies in health care informatics. Int J Biomed Comput. 1994;35(suppl 1):116–27.
27. Information Security Subcommittee, Mayo Clinic/Foundation. Data Security Policies and Standards; September 1994 (provided by Dr. Christopher D. Chute, Section of Medical Information Resources, Mayo Clinic/Foundation, Rochester, MN).
28. Institute of Medicine. Confidentiality and privacy of personal data. In: Donaldson MS, Lohr KN, eds. Health Data in the Information Age: Use, Disclosure, and Privacy. Washington, DC: National Academy Press, 1994.
29. Kowalski S. An accountability server for health care information systems. Int J Biomed Comput. 1994;35(suppl 1):130–8.
30. Kumar S, Spafford EH. An application of pattern matching in intrusion detection. Technical Report CSD-TR-94-013. COAST Project, Dept. of Computer Sciences, Purdue University, West Lafayette, IN, June 17, 1994.
31. Lawrence LM. Safeguarding the confidentiality of automated medical information. Jt Comm J Qual Improv. 1994;20:639–46.
32. New York State CHMIS Executive Policy Committee. NYS CHMIS Confidentiality and Data Security Policy, Draft. Albany, NY, December 21, 1994.
33. Robinson DM. Health information policy: without confidentiality. Int J Biomed Comput. 1994;35(suppl 1):97–104.
34. Shea S. Security versus access: trade-offs are only part of the story. JAMIA. 1994;1:314–5.
35. Shneir B. Applied Cryptography. New York: John Wiley & Sons, 1994.
36. Szolovits P, Kohane I. Against simple universal health-care identifiers. JAMIA. 1994;1:316–9.

1995

37. American Civil Liberties Union of Massachusetts. Statement of opposition to S. 1360, the so-called "Medical Records Confidentiality Act of 1995." Boston, November 4, 1995.
38. Davis R. Online medical records raise privacy fears. USA Today. March 22, 1995:A:1.4.
39. Latham L. Network security, part 2: policy should come first. Inside Gartner Group This Week. April 26, 1995.
40. Safran C, Rind D, Citroen M, Bakker AR, Slack WV, Bleich HL. Protection of confidentiality in the computer-based patient record. MD Comput. 1995;12:187–92.
41. Senate Bill 1360: The Medical Records Confidentiality Act of 1995. 104th Congress, 1st Session.

1996

42. Paskavitz MR, ed. Briefings on the JCAHO. Information management (IM). Special Report Review and Analysis: 1996 JCAHO Standards. Marblehead, MA: Opus Communications, 1995; 16–18 (ISBN 1-885829-14-0).