Medical Records on the Internet

Can Security and Privacy be Assured?

Ron Risley, MS-IV
University of California San Diego, School of Medicine
Dan Masys, MD, committee chair
Chris Mathews, MD, MSPH committee member
Mark H. Ellisman, PHD, committee member
Harold Simon, MD, PHD, advisor

1 May 1996

Abstract

Changing economic conditions are driving rapid change in the organization of health care providers in US metropolitan areas. Mergers, restructuring, and shifting contracts and alliances can create nightmares for those charged with maintaining and distributing patient records. Making records available on the Internet would ease the task of distributing information to and gathering data from new or newly reorganized provider organizations, but it raises serious questions about privacy and security. Despite the common perception that connecting to the Internet is an open invitation to invasion by hackers and industrial spies, advanced technologies offer the promise of convenient, inexpensive, and efficient information distribution that is both more private and more secure than existing methods which rely on paper and proprietary communications networks.

This paper explores techniques available for securely serving information via the Internet and discusses some potential security pitfalls.

Table of Contents

Introduction

To keep your secret is wisdom; but to expect others to keep it is folly.
— Samuel Johnson

Three trends are having an impact on today's medical record keeping environment: managed care, the growth of the Internet and other electronic communications technologies, and an increasing awareness of privacy issues. The growth of managed care organizations has prompted a change in the alliances among health care organizations (Lawrence and Jonas 110). Where once such alliances were fairly stable, and referrals to specialists were often made on the basis of personal and professional acquaintance, now the choice of a specialist is often driven by membership in a specific health plan. Such dynamic provider relationships create problems with distribution of and access to medical records, particularly when record keeping relies on paper-based technology (Weitzman 366). Fortunately, the past few years have seen an explosive growth in electronic information processing technology. Parallel to the growth of information technology, though, is an increasing awareness of and concern about personal privacy. Any implementation of an electronic medical record keeping system must go beyond the standards set by present-day systems to protect the privacy of patients' medical information (Rotenberg).

The Persistence of Paper

Nearly two decades after the "paperless office" became a hot marketing buzzword, medical records are still kept primarily on paper. Even in institutions that support computerized record keeping, the paper record is often still regarded as the first and most accurate source of information for both clinical and medico-legal decision making.

Though the changing face of health care is making paper-based record keeping less practical, less secure, and more costly (Detmer), there are a number of reasons why paper remains the preeminent system for record keeping in spite of dramatic improvements in electronic record keeping technology. Any electronic system which hopes to replace primarily paper-based systems must replicate the strengths while addressing the weaknesses of paper.

The strength of paper based systems lies in their usability without special equipment or training. Any sighted person with the proper linguistic and technical background can read and understand a paper-based record without any special equipment (notwithstanding indecipherable handwriting). Record formats are well standardized. Creating new entries in the record requires only ubiquitous equipment (a pen and paper), and authenticity of records if reasonably simple to verify at a later date using document analysis. Paper can support graphical representation of data as well as text, allowing for emphasis of particularly important points in an intuitive way. The skills required to create, access, store, transport, and maintain a paper record are common in the population and well understood by the many types of personnel—medical, social, clerical, legal, and informatics—who must maintain and manipulate the information in the records.

Indeed, paper-handling technologies are very well developed. While a physician from the 1800's would be amazed and dazzled by MR imagers, cardiopulmonary bypass pumps, broad-spectrum antibiotics, and doppler ultrasound, he would be quite at home with a modern hospital chart.

The Emergence of Electronics

Given the simplicity and accessibility of paper based records, why consider anything else? There are two primary forces driving a switch to newer information-handling technologies. One is the increase in complexity of the health care environment, and the other is the increasing power and decreasing cost of electronic storage and communications systems. Falling prices and increasingly intuitive user interfaces are making electronic systems more accessible at the same time that changes in health care delivery systems are making paper systems less acceptable (Szolovits).

Chief among the problems with paper records is mobility. The original record can exist only in one place. Duplicating the record is expensive, time consuming, and prone to errors and omission. Xerographic copies of the chart are difficult to authenticate, easily viewed and modified by unauthorized persons, and expensive and time-consuming to transport (Kohane, Greenspun, Fackler, Szolovits). When only one authentic copy of a record exists, it is vulnerable to loss from any one of a number of causes. In an era when most global population centers can be reached electronically in seconds, moving paper-based records can take days. The widespread use of facsimile machines has ameliorated this problem somewhat, but faxed records are even more prone to errors, forgeries, and unauthorized access than their paper counterparts.

The Design of the Internet

Over the past three years, one of the most dramatic changes in the landscape of informatics has been the growth in popularity of the Internet . Since the time its progenitor (called Arpanet) was launched by the US Government in 1971 as an experimental communications network designed to survive a nuclear assault, the Internet had largely remained a relatively obscure system used primarily by researchers and students at major universities worldwide. At the same time that computers were becoming more commonplace and technology was providing more powerful processors, faster modems, and cheaper storage devices, political changes in the way the Internet was managed and funded and improvements in the software used for accessing the Internet resulted in explosive growth (Engst 34).

Thanks to its heritage as a decentralized system designed to survive nuclear attack, the Internet is without any central controlling authority. Therefore, the size of the system and the number of users is difficult to ascertain. It is estimated that the Net had over four million users in 1993; some experts believe there may be ten times that number today. Somewhat easier to quantify is the popular awareness of the Internet. In the University of California's Melvyl newspaper database [indexing the *Christian Science Monitor* (National edition), the *New York Times Magazine*, the *Los Angeles Times* (Home edition), the *Wall Street Journal* (Eastern and Western editions), the *New York Times* (Late and National editions), the *New York Times Book Review*, and the *Washington Post* (Final edition)] the word "Internet" appeared as a keyword in only nine articles in 1991. In 1995, the Internet was mentioned in nearly 1,000 articles. This hundred-fold increase in popular awareness has shifted the Internet from the domain of researchers and computer geeks to a ubiquitous feature on America's media landscape where law firms, shopping malls, hospitals, used car dealers, movie promoters, and a sizable number of private individuals all boast home pages on the Internet's World Wide Web.

The Advantage of Ubiquity

The Net seems to be everywhere these days, but does that argue in favor of its use as a server for medical records? I believe that it does. Existing hospital information systems tend to require specific computer systems to operate, they generally cannot interoperate with one another, special training is required to learn to use them, they do not scale well (it is difficult for the same systems to economically serve the needs of both a small group practice and a huge, multi-site, corporate model HMO), and they are relatively inflexible once installed (Kohane *et al.)*.

World Wide Web based information servers, on the other hand, can be accessed from client software that runs on virtually every type of personal computer in common use today, from very inexpensive laptop and palmtop machines to high-end workstations and clusters (Children's Hospital [Boston] W3-ERMS Project). They offer friendly interfaces, and an increasingly large segment of the workforce is already familiar with their use. Web-based systems are easy to decentralize, so that backup copies of critical records can be automatically stored at geographically separated locations (Kohane *et al.)*. They scale well: an individual can easily and economically set up a system that can serve a small operation with an identical interface and seamless interoperation with huge systems serving mega-providers. Web-based systems even share with paper the advantage of being driven by a much larger user base than dedicated hospital information systems: advances in the technology are driven by a large market that demands improvements while at the same time placing a premium on compatibility (Kohane *et al.)*.

The Drawbacks of Accessibility

Being able to instantly access the medical records of a traveller from Iowa who develops an acute abdomen while vacationing in Vienna might be a boon to the traveller and to the medical team, but what about others who might access those records? A significant number of popular press articles on the Internet concern its security weaknesses (Anderson 1996; Borzo; Children's Hospital [Boston] W3-ERMS Project; Detmer; Eid; Kohane *et al.;* Markwell*)*. If medical records are available on the Internet, will that not represent an open door to hackers in dingy basement rooms stealing medical data to sell to unscrupulous employers, government agencies, insurance companies, law enforcement officials, scandal magazines, and even rabid telemarketers? Paper records might have their drawbacks, but at least they require that an attacker be physically present to compromise them.

It is true that the Internet, because of its widespread accessibility, presents remarkable opportunities for information thieves or malicious hackers to wreak havoc. However, it is not true that there are inherent weaknesses in the Internet that make it unsuitable for record storage. Indeed, many systems in use today—both paper and electronic—rely more on complexity and custom to protect information than on any true security. Faxing a discharge summary to a colleague across town may seem to carry little risk, for example. But such an undertaking assumes many things: that the sender is able to accurately identify the recipient, that the telephone company is honest, that the receiving fax machine is attended by secure personnel, that the faxed records will be disposed of properly, that the number to which the records are being faxed is correct, and that the telephone system is itself secure. This last point is particularly important—as telephone systems become more sophisticated they begin to resemble the Internet more than the hard-wired telephone circuits of old. Further, as more and more companies enter long-distance and local service markets, the likelihood for failures and security weaknesses increases dramatically, as do the chances that a given telephone connection will be carried by an easily intercepted radio link. Indeed, the phone system has already been the target of a number of malicious attacks.

Any system which places sensitive information such as medical records on the Internet must be held to a higher standard of security than is currently applied to paper records or proprietary electronic systems. While this places a greater burden on the implementors of an Internet based system, at the same time will result in a system which is far more robust. Instead of relying on security by obscurity, and hoping that an attack will not come, an Internet records server will be designed to withstand attacks. This is crucial, as such attacks are virtually inevitable.

There is one more factor that at least partially ameliorates the burden of providing increased security for an Internet based system. Medical records are not the only pieces of sensitive information being transported on the Net. An active industry has grown up around providing security solutions for the Internet, and many security solutions are currently available off the shelf (Netscape 1995).

These solutions have the advantages of having their development costs spread across a wider customer base than systems dedicated to health care can support, and many of them have been proven or hardened by withstanding—or failing to withstand—innumerable attacks (Kohane, Greenspun, Fackler, Szolovits).

The Ideal Medical Records Server

There are a number of attributes which could be ascribed to the ideal system of disseminating personal medical records. This list is by no means exhaustive, but should serve as a starting point for comparing various approaches. The goal is to construct a framework wherein the proposed Internet based records server can be measured against existing and planned new systems.

Privacy

Privacy is, perhaps, the first issue that confronts anyone who proposes to store and distribute medical records. The concept seems simple, but keeping medical records private is more than just only allowing access to the records by persons authorized by the patient (Nagel).

Levels of Privacy

Different types of medical information demand different levels of privacy. Some information should be readily available with minimal barriers to any medical personnel who request it. For example, the existence of drug allergies or life-threatening chronic conditions such as diabetes mellitus might be critical for an emergency room physician who might otherwise be unable to get permission from the patient, family members, or a primary provider. On the other hand, certain very specialized records (examples might be notes related to acute psychiatric care or participation in a blinded drug trial) need to be kept private, at least temporarily, even from the patient. Between these two extremes are countless combinations: records pertaining to financial issues which payors might require access to, but which the patient might not want the physician to see, information (the results of genetic testing, for example) which might be important to providers but which a patient wishes to keep out of insurance company files, and statistical information which an institution might feel is proprietary and should not be accessible to other providers in the same network but outside the institution.

Access Control

The need for various types of protection for different types of information argues for a system that puts the ultimate authority for access in the hands of the patient. This might seem, at first, to fail to accommodate the scenario described above where the information must be kept even from the patient. Fortunately, this is not necessarily so; an example of how this might work is discussed below (see "Research Data," page 23).

Another problem is that the system seems complex and unwieldy. Is the average consumer of health care services willing to invest the considerable time it might take to fully understand such a system of access permissions and manage it appropriately? The answer is almost certainly "no." The fact that the system might not be necessary for everyone, however, does not argue against making it available to those who desire it. A well-designed secure system could default to a set of access permissions similar to what exists today; that is, control is largely held by the institution creating, managing, and accessing the records. The important point is that any patient could, at any time they deem appropriate, take control of some or all aspects of access to their medical records.

Traffic Analysis

Even if a system manages to keep all records away from prying eyes, it might not fully protect patients' privacy. To be truly private, a system would have to conceal even the existence of certain information, as well as shield the data base from intruders who would count records or monitor changes or retrievals. Otherwise, the system would be vulnerable to traffic analysis (Schneier 219), a technique where information is inferred from the existence or movement of data, even when the data itself is encrypted and therefore unreadable. An insurer, for example, might refuse to provide coverage for an individual who had more than a threshold number of entries in the medical records data base, even when the content of those entries might not be known.

Security

It is not enough that information should be kept from the prying eyes of those who should not have access to it. It is also necessary to insure that information is not lost, destroyed, altered, or fabricated.

It may seem paranoid to think that someone may set out to deliberately destroy or alter medical records, but it pays to remember that much damage is done to records by accident (misfiling, accidental erasure), and a good defense against malicious attacks generally provides a good defense against accidents as well.

Authentication

Security schemes ultimately rely on a chain of accountability. If an individual is held personally accountable for, say, the content of a medical record entry, that person will take certain pains to be sure that such content is reasonably accurate. Thus, we ask persons who contribute to the medical record to sign their entries. Security breaks down in large institutions if personal accountability is not maintained. Imagine a 10,000 person institution where a range of people—file room clerks, physicians, nurses, janitors, couriers, transcriptionists, utilization reviewers, computer programmers, and typists—all have access to a celebrity's records. With so many people to share any blame or suspicion, one would have to be naive to think that the information won't get shared with someone's spouse, friend, or ultimately the press. The ideal records management system, then, would make each person with access to information personally responsible for the security of that information. In order to assign responsibility, though, we have to be able to positively identify those individuals.

Authentication—verifying that a person really is who they claim to be—is a key element in any private and secure records management system (Schneier 52-56). Fortunately the problem of authentication is not unique to the medical industry; much can be learned from the way financial institutions protect money. Still, in the medical world it can be a particularly vexing problem because of the range of individuals who must be covered. While it might be perfectly reasonable to assign a relatively expensive, difficult-to-forge identification such as a smart card to a career employee of a large institution, one must also authenticate the identity of patients, who might for any number of reasons find keeping track of even a credit-card sized piece of identification difficult or impossible. Even if records are kept by individuals using smart cards (credit-card sized secure electronic storage devices), a backup storage system would still be required. Smart cards, therefore, provide no inherent increase in privacy or security (Chaos).

Most authentication schemes employ a multi-part identification system. One part is a token, usually a physical object such as a credit card or ID badge. The token is somehow unique (usually by virtue of being assigned a number), but is generally not too difficult to steal, forge, or duplicate. Another part of the authentication system is a key related to the individual—a signature, photograph, fingerprint, or memorized passphrase such as the personal identification number (PIN) used with automated teller machines. In order to falsify an ID, an attacker must both forge the token and duplicate the key.

Authentication systems accessible to the general public are notoriously unreliable. Few persons really know how to choose an effective passphrase (Spafford 1992), people are likely to divulge their passphrases and loan their tokens when convenience dictates, and people are also prone to losing tokens and forgetting keys. This often is not a problem in a world where most transactions take place face-to-face between people who know each other reasonably well, but as health care provider networks spread out, records become more electronic, and physician-patient relationships become more fluid, relying on face-to-face identification becomes far less effective.

To further complicate matters, any truly strong authentication system, particularly if it relies on universal identifying cards or numbers, might pose a threat to anyone who views a universal identification system as a threat to their personal privacy—often these are the very people most interested in medical records privacy (Detweiler).

The ideal server, then, provides a means to identify every person with access to or control over records in order to be able to hold them personally accountable for those records, while not providing any means to track or otherwise compile a database of information on those individuals. These two goals are somewhat contradictory, yet we will explore some novel ways of at least partially satisfying these two conditions.

Accessibility

While the perfect server would protect the privacy of all information contained on it, it would also allow ready access to information by any *authorized* user anywhere in the world. No system can fully meet both these goals . In a situation analogous to adjusting the sensitivity and specificity of laboratory screening tests, security measures can be made rigorous, excluding nearly all unauthorized users but probably some with legitimate needs as well, or they can be made more lax, insuring accessibility at the risk of allowing some unauthorized access. Given that no system can be perfect, one can at least aim for an ideal system that can be tuned to the needs of individual consumers and specific data types.

In a broader sense, accessibility can also be equated to *availability.* It does not matter how well security measures work if the person with the file room key has gone home for the weekend, or the computer is down, or the chart jams the fax machine. An ideal system would be accessible at all times, from all locales, and would be impervious to acts of God, gremlins, or governments.

Economy

More than ever before, the driving force in health care delivery is economics (Kovner). The industry is adjusting to a world in which prolonging or improving the quality of life are not goals which are to be pursued at any cost, but simply factors on the "benefit" side of a cost-benefit analysis. The ideal medical record keeping system would be adopted not simply because it is A Good Thing, but because it provides real, measurable benefits in terms of reduced cost and more efficient utilization of resources.

Anonymity

Anonymity is an issue closely tied to, but distinctly different from, privacy (Detweiler). If a medical record keeping system existed that was so advanced that it was guaranteed that no unauthorized person or agency would ever have access to private information, there might be no need for anonymity. In the real world, however, there will always be situations where the best guarantee of privacy would be anonymity. As technologies such as genetic testing improve, there will be increasing pressure to be able to provide information to health care consumers that *can not* be readily tied to the individual . While perfect anonymity is difficult to achieve in a health-care setting (if for no other reason than the fact that medical diagnosis and treatment requires that a person by physically present), a good medical record keeping system can and should have features which facilitate anonymous service (Anderson 1996).

Serving Medical Data on the Net

By using some off-the-shelf technologies and a minimum of customizing, a system can be developed that takes advantage of the accessibility of the Internet while taking on many of the attributes of the ideal server.

How the Internet Works

As discussed earlier (see page 4), the Internet physically consists of a number of independent computers communicating with one another over telephone lines or other networks. That communication is governed by *protocols* which define the way information is transferred. The fundamental protocol of the Net is called *TCP/IP*. Layered on top of TCP/IP are protocols to achieve specific functions: several handle electronic mail (*SMTP*, *POP*); some mediate file transfers (*FTP*); others control distribution of Usenet news files (*NNTP*); still others transport interactive plain text (*Telnet*).

In the early 1990s, as personal computers became more powerful and increasing numbers were able to connect to the Internet, demand for a "friendlier" interface drove the development of a new protocol (*hypertext transport protocol*, *HTTP*). HTTP was used as the basic protocol to implement the *World Wide Web*, a collection of interactive documents containing styled text, graphics, and links to and within other documents available on the Net. For the first time[1], it became possible to access the bulk of the Internet without using arcane commands inherited from the Unix operating system.

As with many Internet protocols, HTTP follows a *client/server* model. The server machine runs software which allows client machines to connect and access information. The client machine runs software—called a *browser* on the World Wide Web—which presents the information to the user and allows interaction with it. The terms imply a hierarchy, but it is not necessarily true that

[1] The World Wide Web was actually preceeded by a similar but more limited point and click interface known as Gopher; all current Web browsers are capable of navigating Gopher server sites.

servers are larger or more capable than clients. Often machines serve multiple functions and can simultaneously be a client and a server to several other sites on the Net.

Encryption

Given the description of the Internet, it is easy to see why security can be such a difficult issue. *Any information passing from one machine to another on the Internet it subject to eavesdropping, loss, deliberate destruction, and tampering*(Netscape Communications, Inc.). How, in such an environment, could data ever be expected to be private and secure? One answer is encryption (Office of Technology Assessment).

Encryption is a process of mathematically manipulating information so that it cannot be read by anyone who does not have access to a decryption key, a piece of secret information that decodes the encrypted data. Mathematical encryption has been used since before computers existed to secure information, but a fundamental change took place in the late 1970's. With conventional encryption systems, the same key is used to both encrypt and decrypt data. If Alice wishes to send a secret message to Bob, she first chooses a key, uses that key to encrypt the message, then sends the message to Bob. Even if an eavesdropper (Eve) obtains a copy of the message, she cannot read it without the key. Bob presumably has the key, and can use it to decode the message. The problem, of course, is how Alice gets the key to Bob. She cannot send it to him over the same channel she uses to send the message, because it is likely that Eve would then intercept the key and read the message. Instead the key must be transmitted by a secure channel—during a face-to-face meeting, say, or a code book delivered by a trusted courier (Schneier). These are the things that make spy novels exciting.

Public Key Cryptography

Public key cryptography is significant because it uses *two* keys, or more correctly a single key with two parts. One part of the key is public: if Bob owns the key, he might publish it in directory along with his email address and telephone number. Alice can look up Bob's public key in the directory, and use it to encrypt the message. The public key, however, will not *de*crypt the message—that requires the other part of the key, which Bob has and keeps secret. Once the message is encrypted, only Bob can read it. Now Alice can send a private message to Bob, even if they have never met, and know that only Bob can read it. If Alice also creates a public–private key pair, and publishes the public key, then Bob and Alice can communicate privately with no need to exchange keys via a secure channel (Schneier).

Digital Signatures

The most common public key cryptosystem, known as RSA, can also work in reverse (Schneier). If Bob encrypts a message with his secret key, then it can only be decrypted with his public key. Since his public key is published, what good would this do? It seems as though anyone would be able to read Bob's message.

It is indeed true that anyone could read the message by decrypting it with Bob's public key. However, the fact that Bob's *public* key decrypts the message means that only Bob's *private* key could have encrypted it. It authenticates the message as having come from someone with access to Bob's private key. As long as Bob takes care not to reveal his private key to anyone else, this technique authenticates the original message as having come from Bob: the message has a *digital signature.*

Digital signatures are an important advance in electronic record keeping. Handwritten signatures are only really useful if they appear on original documents. Once a hand-signed document has been faxed, scanned, photocopied, or otherwise reproduced, there is no way to verify with certainty that the copied document was not altered. A digitally signed document, however, cannot be altered without invalidating the signature—copies can be as quickly and accurately authenticated as originals. Recognizing this, the US Government now permits electronic signatures on a variety of medical documents (US Department of Health and Human Services).

A Simple Records Server

Let's see how conventional and public-key encryption might combine to make a simple Internet based medical information server. Imagine that a social worker, Sandy, interviews a patient, Pat. Sandy takes notes during the interview on a personal computer. When the interview concludes, Sandy's private key is used to digitally sign the message, and Pat's public key is used to encrypt it. All this takes place within the confines of Sandy's office; nothing has reached the Internet yet.

Once Sandy has signed the record and it has been encrypted with Pat's public key, it is now safe to publish it on the Internet. Once there, anyone with Internet access can retrieve the record, but no one can read it without Pat's secret key. The record is instantly available almost anywhere on Earth, but cannot actually be read without Pat's authorization.

The Secure Server

The simple scheme addresses one problem, privacy, but it is still a long way from our ideal server. What if Pat should become incapacitated? Unless one were to surrender one's private key in advance, one's records could become lost forever. For that matter, what if Pat merely wants a record sent to a consultant across town. If it is necessary to reveal a private key at each such encounter, the ability to truly authenticate would be lost: there would simply be too many opportunities for someone to steal Pat's private key. Fortunately, protocols can be devised which avoid this pitfall.

Multiple Addressees

In reality, public key cryptosystems are complex and slow. They are seldom used to actually encrypt messages directly. Instead, the message is encrypted using conventional cryptography. A random key is generated, known as the *session key,* at the time the message is created. The message is then encrypted with the session key. The session key, which is presumably much smaller and easier to encrypt than the original message, can then be encrypted with Pat's public key, and the resulting encrypted session key can then be appended to the encrypted message (see Appendix, page 38). Now if Pat is dying to read his medical record while vacationing in Zimbabwe, he need only find a nearby Internet connection, download the encrypted record, use his public key to recover the session key, and use the session key to recover the original record. It sounds pretty complicated, but off-the-shelf software exists which performs all these tasks without the user having to understand the details.

Once the session key is separated from the public key, an opportunity arises for some clever tinkering. For example, we might as well add a copy of the session key encrypted with *Sandy's* public key (see Appendix, Figure A-1, page 40). After all, Sandy wrote the original message, so nothing is lost by allowing her to access it. Nothing is lost, but something is gained. Now Sandy can read the message without Pat having to reveal his private key. If Pat loses or forgets his key or otherwise becomes incapacitated, Sandy can access his records, or at least the ones she wrote. While we are at it, we might as well add another copy of the key, encrypted with Pat's primary care physician's public key. That way Pat's doctor can review the record any time, without needing Pat's key.

Note that all this decrypting and encrypting takes place on private computers in trusted providers' offices. No unencrypted data is carried on the Internet. Almost immediately after the record is created, it can be sealed in a virtual envelope that can only be opened by the patient or those the patient designates. Contrast that with a paper-based system, where everyone who handles the record can conceivably read it. Every copy of a paper record, too, becomes another source for compromise. If the record needs to be used at a remote location, a copy clerk, mailroom personnel, mail handlers, and receptionists will all have potential access to the record. If the record is faxed, it could sit, exposed, on a fax machine until it is picked up. With an encrypted system, copies of the record are every bit as secure as the original.

Time Stamps and Access Tracing

It often is not enough to know that a record was signed; it is often crucial to know *when* a record was created. Paper-based records rely on handwritten dates and insubstantial cues such as the position of pages within a chart to verify timing. Those cues are not available on electronic records.

Keeping Time

One solution is to use a trusted time stamp service. The time stamp service is nothing more than a site on the Internet that accepts messages, appends a date and time to them, and signs them with the private part of a public key. Since the record is already encrypted, the time stamp machine has no knowledge of the information in it. It simply verifies that the record existed at a certain time. The time stamp machine need be trusted only insofar as the organization which operates it must be relied on to not forge time stamps and to keep its private key secret (so others cannot forge time stamps). If there is some doubt about the integrity of a given time stamp service, any number of independent services can be employed. If one becomes compromised, the integrity of the records can still be verify by the other stamps (Schneier).

We have gone a long way toward assuring access: our records are now on the Internet, a global network accessible from almost everywhere, and we have a system whereby the patient, the originator of the record, the patient's doctor, and anyone designated by the patient can access the records. We can go even farther, and set up voting protocols where, say, any five out of a specified group of seven people can access the records if they all cooperate (Schneier). Access is assured, but what have we done to security? As was pointed out on page 10, the more people have access to information, the more likely it is to be anonymously leaked. In the name of accessibility, we have permitted access by the patient, the patient's designees, the designees' designees, and so on. Fortunately, there is a way to keep track of who allowed whom to do what:

Access Tracing

The time stamp and signature can be applied not just to the original record, but to each added addressee as well (Appendix, page 38). That way, if a document is leaked, the time stamp and signature can be used not only to determine which addressee leaked the document, but who authorized that addressee as well. Unfortunately, this system is easily defeated by simply removing the time stamp and signature. For it to be effective, there would have to be general distrust of any records encountered without a genuine time stamp and signature. Even with this weakness, though, it would be an effective deterrent against casually loaning out one's private key, being careless with decrypted records, or indiscriminately designating new addressees.

Indexing and Searching

Now we have records available on the Internet. We can control, on a record-by-record basis, who has access to what. The problem is that the Internet is a vast place, and locating all of a patient's records would be nearly impossible without some type of index. On the other hand, the very presence of such an index might reveal more than the patient wishes to have known by allowing a form of traffic analysis—for example, if a new record of approximately the same size is added to a patient's index each week one might infer that the patient is receiving ongoing physical or psychotherapy; one can draw some pretty good conclusions without having access to the actual data (see "Traffic Analysis," page 9). Once again, though, there are some cryptographic methods that can help.

Zero Knowledge Proofs

Alice claims to have a secret. Bob claims to know that secret. How can Bob prove to Alice that he knows the secret, without revealing the secret to Alice? What if Alice is only *claiming* to know the secret? It might sound like a strange situation, but assume that Alice operates a secure Internet-based medical records server, and Bob is a physician claiming to be treating a third party, Trent. If Bob wishes to gain some information about Trent, he might try to obtain records from Alice claiming to be Trent's physician. Bob would not be able to read the records because they are encrypted, but if Alice either confirms or denies the *existence* of the records, Bob could use that information in a traffic analysis attack (see page 9).

On the other hand, assume that Bob is, in fact, a legitimate provider of health care services to Trent. A malicious interloper, Eve, is listening in on the conversation between Bob and Alice (or perhaps even pretending to be Alice), hoping to get some private information about Trent. Zero-knowledge proofs are mathematical constructions which allow Bob and Alice to prove to each other that they possess some information (in this case, knowledge of the existence of a record), without revealing the information to eavesdroppers or even to each other (Schneier).

Distributed Storage

Even if one could *completely* obscure the identity of persons to whom records belong, there are still ways in which the server site would be vulnerable to traffic-analysis attacks. Since the Internet lacks any fixed, point-to-point connections between specific users, all information on the Net is contained in packets with addressing information. An attacker could monitor all the traffic flowing into and out of a server site and, by matching packet addresses to provider sites, gain some information about the nature of the records.

This type of attack can be made more difficult by not relying on a single server site for storing all information about a client. It is commonplace for file archives on the Internet to be maintained at multiple sites known as mirrors. Mirror sites periodically share information so that all records eventually are maintained at all sites. Information being passed between mirrors does not contain unencrypted information linking individual records to provider sites, so an attacker would have to monitor a potentially large number of servers in order to obtain useful amounts of information from traffic analysis.

Another advantage of using mirror sites is that, by storing information redundantly at geographically diverse locations, they insure that data will be available even in the event the network or an individual server site is disrupted by equipment failure, an attack directed against the site, or a natural disaster; the Net was designed with just these possibilities in mind (see The Design of the Internet, page 4).

There is a down side to using multiple, distributed sites. At larger health-care facilities, it is conceivable that information generated at one location within an institution would be needed by another department within that same institution before it would have a chance to be synchronized with multiple mirror sites. If server sites are picked at random for initial storage of new information, the requesting department would not know where the information was stored. It is a situation likely to arise only in larger institution; an effective solution might be for such institutions to maintain their own server sites. Such sites could be isolated from the Internet as a whole both physically and by network firewalls (see Firewalls, page 32); they would thus be much less sensitive to attack and traffic analysis. They could still mirror with other server sites to assure global access to the records after a short delay (hours to days), without compromising security.

Anonymous Routing

Institutions with more demanding privacy needs could take advantage of another well-developed Internet technology, anonymous re-routing (Detweiler). Originally developed to allow senders of electronic mail to remain anonymous, the technology is now being extended to the World Wide Web.

The principle behind anonymous re-routing is simple. A sender who desires to send a message anonymously picks from a list of anonymous re-routers operated by a number of organizations and individuals. The sender composes the message, adds addressing information to it, then encrypts it using a the public key of the chosen re-router. Since the message—along with its addressing information—is encrypted, it will be unreadable by anyone intercepting it. When it reaches the re-router, it is decrypted and the addressing information is used to forward the message.

There are subtleties associated with re-routing. The security of the system depends on the integrity of the re-router operator; it is customary to send messages through several such (presumably independently operated) sites. Careful traffic analysis can trace messages, requiring that the re-routers delay retransmission for a considerable period of time so that many messages can be re-transmitted in a random order. The resulting necessarily unpredictable time delay could be troublesome for some health care providers. Because they lend themselves so well to nefarious activities, few sites on the Net are willing to host anonymous re-routers, making them scarce and often unreliable. Still, the technology exists as an option when privacy needs are extraordinary.

Research Data

An area poorly addressed by current medical records processing systems is gathering research data for retrospective studies. Common practice is to conduct chart reviews without obtaining special consent from patients, with the understanding that personal information will not be revealed when the study is published (Kaplan and Sadock 2761).

A cryptographically strong medical records storage system would not easily permit such a review. In order to conduct such a survey, a researcher would have to be introduced into each patient's list of permitted viewers. While this might not be of immediate concern to patients, it is in the long-term interests of all users of health care systems to encourage research. To this end, a "research layer" could, with the patient's consent, be added to patient records which provides demographic and diagnostic data with fewer access restrictions than are imposed on personal information. *This necessarily represents a compromise of privacy,* but for most patients the risks would be small.

Implementation and Economics

The advantage to an Internet-based system is wide accessibility. It does not make sense, therefore, to rely on equipment, software, or techniques that are not as widely available is the Internet itself. Until very recently, the only truly commonplace function of the Net was electronic mail. Since 1993, however, the phenomenal growth of the World Wide Web has greatly enriched the range of services commonly available—Web browser software exists for every computer platform in common use in the world today.

The growth of the Web and its exploitation for commercial purposes has also led to a proliferation of security technologies, primarily designed to protect electronic commercial transactions, allowing credit card data to be transferred without risk of interception. At first glance, these technologies would seem to be well suited for protecting medical data as well.

Unfortunately, there are significant ways in which current World Wide Web servers and browsers fall short of being able to implement a secure medical information server. Most important is that current client-server models concern themselves with protecting the link between a cleartext client machine and a cleartext server data base. The key to the secure medical information server is that data is clear *only* on the client machine, *never* on the server. While current servers and browsers do not directly support this *client to client encryption* model, both can be extended with relative ease to accommodate the needs of the secure medical information server.

Servers

Surprisingly, there are few security considerations for the server side software. The reason is that the secure medical information server architecture makes the assumption that the server is not secure. Without this assumption, every site from which information is served (recall that, to protect against server and network failures, all records are to be stored at multiple sites) must be staffed by trusted personnel at any point where such personnel might gain access to the information. This would also violate the requirement that we maintain strict *individual* responsibility for control of the information.

What the servers must do is provide rapid access to large numbers of binary records; this requirement is by no means unique to medical applications, and many systems exist which can fulfill this need. Servers should also be able to provide synchronization—multiple sites should be able to exchange information so that each holds as complete and up-to-date a collection of data as is possible.

Servers could provide other services as well, such as time stamping and anonymous re-routing, but these are also off-the-shelf functions.

Who Pays?

Servers might be available off-the-shelf, and thus be accessible and inexpensive, but they are not free, nor are they trivial to set up and maintain. The Net is in an uneasy transition from a time when most servers were installed and maintained to fulfill a specific need (often government funded), and excess capacity was given over for free use by users of the Net without a great deal of regard for cost accounting. The past few years have seen a rapid pace of commercialization of the Net, putting increasing demands on the Internet's infrastructure. There are ongoing debates about how best to fund services which are provided by a few but which benefit all. Some striking parallels can be drawn between this development and the corporatization of medicine in the US.

Ideally, large health care institutions would find that the benefits they derive from maintaining open servers (cost savings come in the form of being able to use remote sites for back up) more than offsets the cost of allowing others to store data on their servers. Whether this will actually turn out to be true, and whether cost accountants will recognize the fact, remains to be seen.

Other methods have been proposed for paying for distributed database services. One system proposed by Digital Equipment Corporation relies on charging very small amounts of money, on the order of thousandths of a cent, for each transaction in a barter-type arrangement, and only actually computing payments when minimum amounts have accumulated.

Clients

The client side is where the action is. No off-the-shelf World Wide Web browsers currently implement the sort of client-to-client encryption proposed for the secure medical information server. Fortunately, most browsers support the use of "helper applications," programs which execute automatically when the browser detects that it is accessing a specified file type. It is possible to combine that capability with the scripting services available on Unix, Macintosh, and Windows-based platforms to link the World Wide Web browser to PGP, a *de facto* standard public-key encryption program. Such a system would provide, at very low-cost, a platform with the following attributes:

- Able to be implemented on all commonly-used platforms with largely off-the-shelf, mass-market components, assuring low cost, vendor independence, and platform independence.
- All components of the system critical to encryption and decryption are available in source-code format, allowing independent auditing and verification of security[1]. This is a critical feature often unavailable with proprietary software (Schneier; Risley).
- Does not require a proprietary operating environment; in many cases can exist side-by-side on the same machine as existing systems.

There are some aspects of the system which would be desirable in sites with more stringent privacy requirements that would require additional components (anonymous re-routing, use of randomly selected servers and timestamp services are examples). Users with these additional requirements would be somewhat more limited in their choice of platforms.

Also, while the proposed methodology is compatible with all modern operating environments, a surprising number of hospitals continue to employ systems based on a central mainframe/dumb terminal model popular in the 1970's. While World Wide Web browser software has been successfully implemented in such environments (Children's Hospital [Boston] W3-ERMS Project), public-key encryption techniques generally are much less secure when the hardware performing the encryption is shared by many users (Zimmerman). This will become less of an issue as hospital information systems modernize.

[1] Currently, versions of PGP distributed in the United States and Canada use some routines which are proprietary to RSA Data Security, Inc. and are not available in source form to the general public. This is likely to continue to be the case until the RSA patent expires on 20 September 2000. Versions of PGP distributed outside the US and Canada do not have this restriction.

Communications

Any attempt at a global, distributed server for medical records must acknowledge that health care providers operate in a wide variety of settings, from mobile practitioners in rural areas to one- or two-person offices to large multi-campus hospitals (Kovner). An Internet-based system can reach all these environments, though at present radio-based mobile Internet connections are generally expensive and unreliable. The majority of personal and small business computer systems already have the necessary hardware and software to connect to the Net; most of the rest can be upgraded to be Net-capable for less than $200. Connecting individual small computer systems to the Internet is relatively simple, inexpensive, and vendor-independent (Engst).

Parallel Implementation

Larger institutions are generally already equipped with some type of networking architecture that connects their computer systems. Often the investment in network hardware and wiring exceeds the cost of the computers themselves. For these institutions, the cost of changing data delivery systems can be substantial. An Internet-based system can provide a substantial cost savings in these cases.

We have already seen that the resources necessary to implement the secure medical information server are already present in most computer systems and that a World Wide Web-based system can operate side-by-side with existing systems on the same machine (page 26). Often an institution can begin by equipping only a subset of machines in critical areas with Internet access via a dial-up modem at minimal cost. As use of the new system expands, individual dial-up connections can often be replaced by accessing the Net across existing networks; such gateways to allow TCP/IP (Internet) communication across Ethernet, LocalTalk, and token-ring networks are in common use today.

Even when integrated with a local network, having a few stations equipped with modems provides a simple way to protect against the a disaster should the local network fail. If records are redundantly distributed across the global Internet, modem connections could bypass the equipment failures.

Legacy Systems

Institutions using legacy systems based on a central mainframe and remote terminal model might be more challenged. Since many such systems support some form of electronic mail, and in many cases gateways exist to allow email access to the Internet, some form of parallel implementation might still be possible.

Many users of legacy systems have discovered that, as the dumb terminals that make up the points of access to their system break down, the most cost-effective repair is to replace them with low-end personal computer systems running terminal emulation software. Thus, a good many sites already have hardware in place capable of accessing the Net. Since the MIME protocol we called upon to implement our server over the World Wide Web works equally well for email, these systems' email links could be pressed into service to allow access to Web-based services without additional hardware or server end software. What would be required, however, is software at the client site to move email data into and out of the terminal emulation program. Such software, while not technically difficult to write, would nonetheless represent an investment of engineering dollars for the sole purpose of maintaining an obsolete system. Whether that makes economic sense would need to be evaluated on a case-by-case basis.

Potential Pitfalls

Even after records are encrypted, access is controlled, steps are taken to reduce the opportunities for traffic analysis, documents are electronically signed, and document tracking is implemented, there are still opportunities for mischief. Attacks can take place where information is created, where it is stored, and where it is used.

Though the techniques outlined can be used to create a system in which security and privacy are easier to achieve than with current paper-based and proprietary electronic record keeping systems, no technique can succeed without the active involvement of the participants. In turn, it is important for the users of the system to understand the ways in which it might be compromised.

Server Attacks

Because all information on the servers is encrypted, privacy would not be compromised even if a hacker should succeed in gaining access to a server's contents. Because all individual records in the system are signed using strong cryptographic techniques, attempts to compromise the system by placing forged information in it would not be successful. Indeed, the data—always in its encrypted form—could be made freely available to anyone without risk of compromise. This does not mean, unfortunately, that the system as a whole would be invulnerable to attacks directed at the server. We have already seen, for example, how traffic analysis can be used to gain information even from encrypted records (see page 9). There are still other types of attack on the server network:

Denial of Service

While it might be difficult to access data by attacking the server, it is possible that a malicious attacker (or technological failure) could render the data inaccessible. In addition to allowing an individual to simply wreak havoc on the health care delivery system, the potential for broad-based denial of service attacks could open the door for an attack to hold the secure medical information server hostage by threatening to disable it in exchange for some consideration. Denial of service attacks could also indirectly compromise privacy by forcing providers to go outside the system to equip themselves with potentially insecure methods of backup information storage and retrieval to use when the primary system fails.

In general, the threat of denial of service attacks is limited in the proposed Internet-based system because data can so easily be stored at multiple server sites. Large institutions would most likely have at least one such server for their own patient data located behind a firewall, increasing its resistance to attack (Ranum, Leibowitz, Chapman, Boyle).

Even with distributed storage, however, the Internet is subject to certain network-wide attacks. One such attack was the infamous “Morris worm,” a program which exploited certain weaknesses in common Internet servers to spread itself throughout the Net and thereby adversely affect performance (Spafford 1988). Though the precise conditions which made that attack possible no longer exist, it remains a reminder of the Net’s vulnerabilities (as well as an illustration of how quickly the Internet community can respond even to sophisticated attacks). Other potential weaknesses in the Net have been identified; for example, the technique used to route messages through the constantly-changing architecture of the Net has failed on several occasions (Wallich).

As the Net increasingly becomes a ubiquitous vehicle for commerce and entertainment, it will grow as a target for malicious attack. Ironically, this might well serve to make it *more* suitable for distribution of medical data. Attacks on the Net as a whole will more likely be directed at commercial sites, and the industry at large will have incentive to anticipate, prevent, and limit such attacks.

Keyserver Attacks

Systems based on public-key encryption generally require that there be a correspondence maintained between keys and their holders. In a large, distributed system, public keys are generally retrieved from a directory known as a keyserver. If someone were able to compromise the function of the keyserver, they would be able to insert erroneous information into the medical data base. Normally, forged information could be detected because the signature would not match with the signer’s public key. If an attacker generated a new key, used it to sign false information, then replaced the legitimate person’s public key listing on the keyserver with their own, the signature would appear to match.

Unless a recipient became aware of the key substitution, the data would be regarded as valid.

Fortunately, keyserver attacks are not unique to the medical records system, and several techniques have evolved to thwart them (Schneier 178). The keys are public, after all, and can be widely published, so that a potential attacker would have to replace the key in all locations that users are likely to check. Public keys can, themselves, be signed by one or more central key signing authorities whose keys would be so widely known as to make substitution impossible (RSA Data Security, Inc.). Another approach is used by PGP, which creates a "web of trust" by allowing keys to be signed by any number of other users; the goal is to be able to verify any key by following a chain of signers back to one whom you know by personal experience to be valid (Zimmerman).

Client Attacks

Client-to-client encryption provides excellent privacy and good security against attacks directed at the server. Data at the client machine, however, exists as plain text. Further, the client machine must have access to users *private* keys in order to encrypt, sign, and decrypt messages. Because of the very nature of the distributed systems, client machines will exist in many configurations and will not be under the control of a central agency, as is often the case with traditional electronic medical records systems.

While it may seem that client-side considerations represent an intractable security hole, in fact these problems would likely exist with *any* system which allows remote access to medical data. Nonetheless, awareness of possible attacks is key to their prevention.

Careless Configuration

The biggest threat comes not from malicious hackers, but from careless or improper configuration of client machines. Client machines must be free of software which can inadvertently act as file servers which would potentially make plaintext records available to anyone on the Net. Ideally, any machine being used as a client manipulating secure data would have *no* server software installed or enabled; in general, installation and configuration of server software should be left to professionals who understand the often confusing and obscure subtleties of file protection.

A more subtle threat comes from innocently-installed software that inadvertently creates a security hole. Several popular system extensions available under Windows and Macintosh OS environments, for example, record every keystroke typed in order to facilitate recovery of data in the event of a system crash. Unfortunately, such programs would also allow the recovery of confidential patient information, and even the passphrases used to protect private keys (Zimmerman). An unscrupulous user of the system could even install such software with malicious intent, with little risk of discovery and virtually no risk of prosecution if discovered.

Viruses and Trojan Horses

Even if client machines are properly configured and kept physically secure, the very mechanisms which enable access to remotely-stored medical records also allow access to a myriad of games and utility programs available for download on the Net. An attacker could easily code a virus (a program which, when run, replicates itself into the operating system and/or other programs on the same machine) or a Trojan horse (a program which surreptitiously performs some function other than what the user has been led to expect) which would make confidential data available to the attacker over the Net. Though strict policies against the use of downloaded software might help, such policies have been notoriously ineffective. In addition, recent extensions to Web browser software allow a server to remotely execute programs on the client without user authorization. Though steps have been taken to make this facility resistant to use as a Trojan horse, it has already been exploited to this end (Sun Microsystems).

Firewalls

One defense against client-side attacks from the Net is installation of a firewall. A firewall is a dedicated machine that is inserted logically between a network of client machines and the Internet (Ranum, Leibowitz, Chapman, Boyle). It monitors traffic between the clients and the Net, allowing only a certain well-defined subset message types to pass between the clients and the rest of the world. Firewalls can be effective, but are by no means a complete solution. In addition, installation and maintenance of a firewall is not trivial, and makes sense only for sites which have a significant number of client machines running on a local network.

The Future

It appears inevitable that the Internet and other information delivery technologies will continue to grow more powerful, and there will be continued pressure on health care providers to provide efficient care in a rapidly changing environment while preserving patients' privacy. These trends will drive the creation of new ways of storing and delivering patient data. With some foresight and planning, a smooth transition can be made avoiding both the chaos of multiple, incompatible systems (Children's Hospital [Boston] W3-ERMS Project) or the danger of lax attention to privacy (Office of Technology Assessment).

Other trends, too, will impact the shape of future medical records systems. Awareness of these trends while designing records systems now will help prevent difficulties in the future.

Personal Devices

Personal Digital Assistants (PDAs) are hand held computers designed to perform the functions of paper-based personal organizers and references. Some, such as Apple Computer, Inc.'s Newton MessagePad, are finding wide acceptance within the medical community (Risley 1995). For many physicians they are already replacing traditional patient tracking forms and medical reference texts. Studies at Brigham and Women's Hospital and Wright-Patterson AFB Hospital (Ebell, Dake) have shown that housestaff operate more efficiently when using PDAs, and have demonstrated that PDAs can efficiently replace not just pocket reference books, but online terminals for accessing patient data as well.

Newer generation PDAs are able to access the World Wide Web through telephone connections and wireless links. Implementing a Web-based records server system now will enable a smooth transition to use of personal devices as they become commonplace.

Electronic References

The pace of medical research continues to increase; physicians commonly decry the enormous investment of time required to stay abreast of even a tiny subset of new clinically-relevant information. Increasingly, practitioners are taking advantage of online electronic reference services to provide up-to-date information on demand. Many of these service are already Web based; those that are not can generally be accessed via a Web gateway. The demand to provide Web access at point of care will increase as care standards change to reflect this new reality. Providing access to patient records through the same channels—provided that it can be done securely—will lower costs and improve efficiency.

Telemedicine

Pressure to hold down costs is driving a shift away from specialization, though the US has an uneven distribution of specialists which results in shortages in many rural and urban areas. One alternative for reaching these underserved populations is telemedicine; the proliferation of inexpensive Internet connections and telecommunications equipment will serve to make telemedicine more attractive. Increased use of telemedicine, however, will increase the need for rapid and efficient transport of patient data. Unfortunately, if tools and protocols are not in place to accomplish such transport, the temptation to send patient data via unsecured email, telephone, or fax will only increase. Implementation of compatible, widely-available secure systems now is the best prophylaxis against the erosion of privacy that can come from shifts to modern telecommunications systems (Nagel).

• • •

Clearly there is a trend toward increased use of electronic data processing systems to store and distribute medical records. Though many perceive this trend toward mechanization as a threat to privacy, proper application of cryptographic techniques will serve to increase, rather than compromise, patients' control over their personal medical information.

Resources and Works Cited

Anderson, Ross J. *Patient Confidentiality At Risk from NHS Wide Networking.* 1996.
ftp://ftp.cl.cam.ac.uk/users/rja14/hcs96.ps.Z

Anderson, Ross J. *Security In Clinical Information Systems.* Computer Laboratory University of Cambridge. 1996
http://www.cl.cam.ac.uk/users/rja14/policy11/policy11.html

Anderson, Ross. "NHS Wide Networking and Patient Confidentiality: Britain seems headed for a poor solution." *BMJ* volume 311 no 6996. 1 July 1995 pp 5-6.
ftp://ftp.cl.cam.ac.uk/users/rja14/bmj.ps.Z

Bach, Eric Bellovin, Steve, Bernstein, Dan; Bolyard, Nelson; Ellison, Carl; Gillogly, Jim; Gleason, Mike; Gwyn, Doug; O'Connor, Luke; Patti, Tony; Setzer, William. *Cryptography FAQ: version 1.0.* 1994.
ftp://rpub.cl.msu.edu//pub/crypt/sci.crypt/sci.crypt-faq.txt

Borzo, Greg. "Electronic Data Standards Group Set to Accept Members." *American Medical News.* April 3, 1995 pg 6.

Chaos, David. "Smart Cards in Medicine—Social Aspects." Usenet correspondence, comp.org.cpsr.talk. 13 April 1994.

Children's Hospital [Boston] W3-ERMS Project. *MedWeb Project: Specific Aims and Significance.* 1995.
http://www.chip.org:80/chip/projects/w3emrs/w3emrs.html.

Detmer, William M. *WWW and the Electonic Medical Record.* 29 November 1994.
http://www-camis.stanford.edu:80/people/bdetmer/WWWTalk/WWW-outline.html

Detweiler, L. *Anonymity on the Internet.* 9 May 1993.
http://www.cis.ohio-state.edu/hypertext/faq/usenet/net-anonymity/top.html

Ebell, Hale; Buchanan, Dake. Handheld Computers for Family Physicians. Appleton & Lange: 1995.
http://www.phymac.med.wayne.edu/jfp/art1095a.htm

Eid, Troy A. *Roadblocks on the Information Superhighway: Removing the Legal & Policy Barriers to Telemedicine.* Presented at the *Society and the Future of Computing* conference, Durango, CO 1995.
http://www.lanl.gov/SFC/95/sfcHome.html

Engst, Adam C. *The Internet Starter Kit for Macintosh.*Hayden Books. Indianapolis: 1993.

Equifax, Inc. Computerized Patient Records Seen as Beneficial Trend. 31 October 1995.
http://www.equifax.com:80/headline/octob95/health1.

Equifax, Inc. Equifax Health Information Privacy Principles. 1995.
http://www.equifax.com:80/consumer/docs/health.html

Hripcsak, George; Sideli, Robert. *Introduction to Medical Informatics: Online Lecture Notes.* 1995. http://www.cpmc.columbia.edu/edu/textbook/

Kaplan, Harold I.; Sadock, Benjamin J. eds. *Comprehensive Textbook of Psychiatry VI.* Sixth edition. Williams & Wilkins. Baltimore, 1995.

Kohane, Isaac S.; Greenspun, Philip; Fackler,James; Szolovits, Peter. *Accessing Pediatric Electronic Medical Record Systems Via The World Wide Web.* 1995. http://www.chip.org:80/chip/projects/w3emrs/emrsviawww.html. 1995.

Kohane, Isaac S.; van Wingerde, F.J.; Fackler, James; Szolovits, Peter. The MedWeb Project. http://www.chip.org:80/chip/projects/w3emrs/w3emrs.html

Kovner, Anthony R., ed. *Health Care Delivery in the United States,* fourth edition. Springer Publishing Company. New York: 1990.

Krsul, Ivan. *Authorship Analysis: Identifying The Author of a Program.* Technical Report CSD-TR-94-030. The COAST Project, Department of Computer Sciences, Purdue University 1994. ftp://coast.cs.purdue.edu/pub/COAST/papers/krsul-authorship_analysis_NISSC.ps.Z

Lawrence, Robert S. and Jonas, Steven. "Ambulatory Care." *Health Care Delivery in the United States,* fourth edition Kovner, Anthony R., ed. New York: Springer Publishing Company 1990 pp 106-140.

Markwell, DC. "Fear of Flowing." *Extract from The Proceedings of the 1995 Annual Conference of The Primary Health Care Specialist Group of the British Computer Society* (Electronic Production by Rob Wilson & Sheila Teasdale). 1995 http://www.ncl.ac.uk/~nphcare/PHCSG/conference/camb95/reading.htm

Nagel, Denise. "Medical Privacy? Technology is Eroding Confidentiality." *Newsday* 27 October 1995. http://www.epic.org/privacy/medical/nagel.txt

Netscape Communications Corporation. *Netscape SSL.* 1995. http://home.netscape.com/newsref/std/SSL.html

Novarro, Leonard. "Company's Aim: Put Hospital Paperwork on Terminal List." *The San Diego Union-Tribune.* May 30, 1995 pg C-6.

Office of Technology Assessment. *Protecting Privacy in Computerized Medical Information.* 1994. ftp://ftp.cl.cam.ac.uk/users/rja14/ota.ps.Z

Patrikas, Elaine O. "Canadian Health Care Afflications for the Internet." Usenet: sci.med.telemedicine. 15 Dec 1993.

Ranum, Marcus; Leibowitz, Allen; Chapman, Brent; Boyle, Brian. *Firewalls FAQ Rev 4.* 1995. http://www.tis.com/Home/Personal/Ranum/Page.html.

Risley, Ron. "Keeping secrets by telling it all." *MacWEEK* v2, n3 (Jan 19, 1988):72.

Risley, Ron. "The Medical Students' Guide to the Newton." 1995. http://sdcc3.ucsd.edu/~rrisley/msgn.html

Rotenberg, Mark. "British Doctors Boycott Medical Network." *EPIC Alert 2.13*. Electronic Privacy Information Center (EPIC) October, 1995. http://www.epic.org/alert/EPIC_Alert_2.13.txt

Rotenberg, Mark. "Designing a Good Medical Privacy Bill." *EPIC Alert 2.13*. Electronic Privacy Information Center (EPIC) October, 1995. http://www.epic.org/alert/EPIC_Alert_2.13.txt

RSA Data Security, Inc. RSA's Frequently Asked Questions About Today's Cryptography. 1996. http://www.rsa.com/rsalabs/faq/faq_km.html

Sandlin, Nina. "The Internet: What's on it for you?" *American Medical News.* June 12, 1995 pg 10-12.

Schneier, Bruce. *Applied Cryptography Second Edition: protocols, algorithms, and source code in C.* John Wiley & Sons, Inc. New York 1996.

Spafford, Eugene H. "OPUS: Preventing Weak Password Choices." *Computers & Security* 11(3), pp 273-278. May 1992.

Spafford, Eugene H. "The Internet Worm Program: An Analysis." Purdue Technical Report CSD-TR-823. 1988. ftp://coast.cs.purdue.edu/pub/doc/morris_worm/spaf-IWorm-paper-CCR.ps.Z

Sun Microsystems, Inc. *Java Security FAQ: version 1.0.* 1996. http://java.sun.com/sfaq/.

Szolovits, Peter. "A Revolution in Electronic Medical Record Systems via the World Wide Web." An informal extended abstract prepared for a talk at the conference *The Use of Internet and World-Wide Web for Telematics in Healthcare*, sponsored by the International Association for the Advance of Health Information Technology (IAHIT), Geneva, Switzerland, September 6-8, 1995. http://www.emrs.org:80/medweb/publications/IAHIT.html.

Szolovits, Peter; Kohane, Isaac. *Against Simple Universal Health-Care Identifiers.* 1995. http://www.chip.org:80/chip/projects/w3emrs/againstsimpleidents.html.

United States Department of Defense. *DoD 5200.28-STD Department of Defense Standard: Department Of Trusted Computer System Evaluation Criteria*. 1985. http://all.net:80/books/orange/top.html

US Department of Health and Human Services, Food and Drug Administration. "21 CFR Part 11: Electronic Signatures; Electronic Records; Proposed Rule." *Federal Register.* US Government Printing Office. 31 August 1994.

Wallich, Paul. "A Rogue's Routing: Hackers may ignore individual PC's and undermine the Net." *Scientific American.* May 1995 pg 31.

Weitzman, Beth C. "The Quality of Care: Assessment and Assurance." *Health Care Delivery in the United States,* fourth edition Kovner, Anthony R., ed. New York: Springer Publishing Company 1990 pp 353-380.

Woodward, B. "The Computer-Based Patient Record and Confidentiality." *The New England Journal of Medicine.* 1419+. 333.21 November 23 1995.

Zimmerman, Philip. *PGP User's Guide.* 11 October 1994. http://web.mit.edu/network/pgp

Appendix:
Introducer Records: A Cryptographic Technique for Tracking Authenticated Information

Introduction

The more people who know a secret, the more likely it is to be compromised. This is particularly so in the realm of electronic communication, where those who leak secrets can easily remain anonymous.

Electronic information is also easily forged, and various cryptographic techniques, such as digital signatures, have been developed to aid in the authentication of sensitive information. Unfortunately, such authenticated documents can be even more damaging when leaked, as their authorship can easily be confirmed. When large numbers of people must have access to authenticated information, as is the case with electronic medical records, it is virtually impossible to assure secrecy.

Described here is a technique which separates authentication information from an original document, and attaches it instead to distribution information. When combined with an end to end public key encryption technique, it becomes impossible to disclose authenticated information without also disclosing the identity of the discloser (though disclosing the original document without authentication can still be done anonymously). By restoring individual accountability, this technique could be used to improve the level of privacy which can be achieved over large-scale distribution systems for sensitive documents.

Method

Authentication of electronic documents—verifying their integrity after storage or transmission by insecure equipment—can be achieved using electronic signatures. Electronic signatures rely on the characteristic symmetry of many public-key encryption systems: messages encrypted with a private key can be successfully decrypted with the corresponding public key, thus verifying that the message was encrypted by someone with access to the private key (presumably only the key's owner).

Conventional Digital Signatures

Public key cryptosystems are computationally complex, so in practical systems only a *digest* of the original message is used to generate the signature. A digest is a relatively short number mathematically generated from the original message in such a way that it would be computationally unfeasable to generate another message with the same digest. The sequence for sending such a digitally signed message is as follows:

- Signer computes digest of the original message.
- Signer encrypts digest with signer's private key.
- Signer transmits original message and encrypted digest to recipient.

The recipient can verify the authenticity of the message as follows:

- Recipient computes the digest of the original message.
- Recipient obtains a certified copy of signer's public key.
- Recipient decrypts the encrypted digest received from the signer.
- If the digests do not match, the document has been altered or forged.

Conventional Signatures with Encryption

Digital signatures are often combined with public-key encryption to create messages which are both private and authenticated. Again, because public key encryption is computationally intensive, a shortcut is usually employed. Typically, messages are encrypted using conventional secret-key cryptographic techniques. The secret key, known as a *session key*, is chosen randomly. Like a message digest, the session key is generally much shorter than the original message and can, therefore, be encrypted more quickly. The session key is then encrypted using the public key of the intended recipient. The recipient uses their corresponding private key to decrypt the session key, and uses the session key to decrypt the original message.

The technique of using a separate session key also simplifies the process of sending the same message securely to multiple recipients. Instead of re-encrypting the entire message for each recipient, the message can be encrypted once, and records can be appended that contain copies of the session key encrypted for each recipient.

When combined with digital signature techniques, messages typically consist of three seperable parts: the original message encrypted with the session key, a collection of session keys encrypted with each intended recipient's public key, and the digest of the original message encrypted with the signer's private key (see Figure A-1).

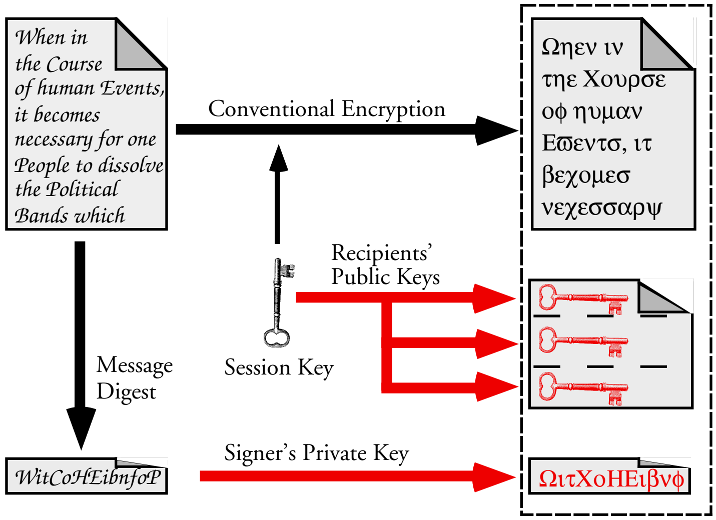
When in the Course of human Events, it becomes necessary for one People to dissolve the Political Bands which
Conventional Encryption
Ωηεν ιν τηε Χουρσε οφ ηυμαν Εϖεντσ, ιτ βεχομεσ νεχεσσαρψ
Recipients' Public Keys
Message Digest
Session Key
Signer's Private Key
WitCoHEibnfoP
ΩιτΧοΗΕιβνφ

Figure A-1: Typical public key encryption/authentication message generation. Public key encryption operations and results are shown in color.

A drawback to this approach is that it requires that all recipients be completely trusted with the information. All recipients are free to pass the information on to others securely (by appending the session key encrypted by the new recipient's public key) or to divulge the information to the public. The information thus revealed can be authenticated using the electronic signature, and there is no personal accountability: any person granted access to the original message, either directly by its originator or as an nth generation recipient, shares equally in responsibility for leaks.

Accountability

When information is likely to be shared among dozens or hundreds of recipient (as is often the case with medical records) it is unreasonable to expect them to be kept private unless recipients remain personally accountable for the dissemination of that information. Fortunately, when information must be both encrypted and authenticated, the authentication information can be bundled with the distribution information. Suppose Alice wishes to send some sensitive information about Duane's substance abuse treatment to Bob, a social worker, and Carol, an insurance adjuster. Alice is concerned that Bob or Carol might divulge the information to Duane's employer, either directly or by passing the information to others who might divulge it. Neither Bob nor Carol have any particular incentive to keep the information confidential, as they have no direct connection to Duane nor could they be held directly accountable should the information appear in the hands of Duane's employer. The only place to point fingers would be at Alice, whose electronic signature appears on the record.

Suppose that, instead of signing the document prior to distribution, Alice instead posts the document encrypted with the session key but without an electronic signature. When Bob requests access to the document, Alice prepares a special record called an *introducer record* which contains the session key encrypted with Bob's public key (exactly as in the conventional example), plus a string of text that positively identifies Bob (such as "This message was prepared especially for Bob") plus an electronic signature consisting of an encrypted digest of both the original message *and* the identifier string (see Figure A-2). Bob can still verify that Alice signed the original message, but he can only check the signature *after* appending "This message was prepared especially for Bob." (In practice, the identifying string would probably be a digest of Bob's public key.)

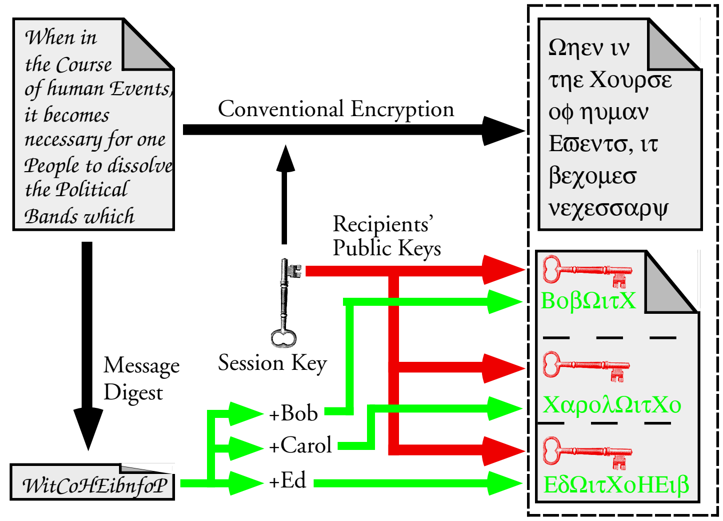
When in the Course of human Events, it becomes necessary for one People to dissolve the Political Bands which
Conventional Encryption
Ωηεν ιν τηε Χουρσε οφ ηυμαν Εϖεντσ, ιτ βεχομεσ νεχεσσαρψ
Recipients' Public Keys
Session Key
Message Digest
WitCoHEibnfoP
+Bob
+Carol
+Ed
BoβΩιτX
ΧαρολΩιτΧο
EδΩιτXoHEιβ

Figure A-2: Introducers make the identity of recipients an unalterable part of the document signature. Public key encryption operations and results are shown in red. Signing operations (using the signer's private key) are shown in green.

Bob could still decrypt the original message and surreptitiously distribute it to others, but only without Alice's signature. There would be no way for recipients of the document to be sure that Bob (or someone else) did not forge it. If Bob needs to give the record to others for a legitimate purpose, he can do so by sending them an introducer record. No matter how many times an authenticated document gets redistributed, a path can always be traced from the original signer to the last person to legitimately have access to the record.

Application

Clearly, introducer records do not provide absolute security. Their strength lies in the fact that distribution information is tightly linked to authentication information. In an environment where authentication was largely inherent in the documentation itself (if the electronic record were scanned copies of signed handwritten notes, for example) there might be little lost by removing electronic signatures. As fully electronic records become more commonplace, however, electronic signatures will become more important. Ironically, the ease with which (unsigned) electronic records can be easily forged will contribute to the enforcement of accountability through the use of a technology such as introducer records.

Colophon

This body text of this document is set in Adobe Garamond, using the Adobe Garamond Expert face for old-style figures, ligatures, and small caps. Titles and headings are set in Avant Garde.

The text and graphics were prepared on an Apple Macintosh PowerBook 5300cs using Nisus Writer software for text and Canvas for graphics. Conversion to Portable Document Format was performed by Adobe Acrobat Distiller.