Privacy and Security in Electronic Health

Dr. Stefan Brands
Credentica Inc.
brands@credentica.com

Version 1.0 of March 10, 2003*

Abstract: This paper describes the growing privacy and security needs in electronic health, and explains why they cannot be addressed by today's commercially available security technologies. A solution is proposed that allows electronic health information to be jointly managed by the data subject and his health care professionals, in a manner that simultaneously protects the data subject's privacy interests, the professionals' liability interests, and the legitimate interests of researchers and other third parties.

"*As more of our medical records are stored electronically, the threats to all our privacy increase.*"
"State of the Union" Address of William J. Clinton, United States, January 19, 1999.

1. Electronic Health Records

Paper-based health records are rapidly becoming outdated. They are easily lost, are subject to wear and tear, are costly to handle, cannot be transferred electronically, may be difficult to interpret, and are inefficient. These fundamental drawbacks are driving a transition across the globe towards Electronic Health Records (EHRs). The Office of Health and the Information Highway (Health Canada's focal point for the use of information and communications technologies in the health sector) [1] defines an EHR as "The health record of an individual that is accessible online from many separate, interoperable automated systems within an electronic network." EHRs can contain a variety of data and be used for different purposes by different parties involved in health care. A literature search reveals an abundance of related terms, including Computer-based Patient Records, Patient-carried Patient Records, Computerized Medical Records, Electronic Patient Records, Electronic Medical Records, Personal Health Records, Integrated Care Records, and Card-Based Patient Records. EHR has become most widely used term, encompassing all reasonable notions of health records in electronic form.

Currently, the vast majority of health records are still in non-electronic form, dispersed across many institutions. However, the ever-increasing popularity of the Internet and electronic communication devices (such as Web-enabled cell phones and PDAs) is causing a rapid shift. According to a survey in 2001 by Harris Interactive and ARiA Marketing, patients are increasingly interested in communicating with their doctors, obtaining personalized medical alerts specific to their medical histories from their doctors, and participating in on-line communities to get information on diseases, hear about alternative therapies, and chat with others. On the side of health service providers, major efforts in Web connectivity are underway. The popularity of handheld devices for medical purposes is rising fast, especially amongst physicians in Europe; according to a report by Harris Interactive in August 2002, 11% of practitioners in the EU use PDAs in their practice, the Netherlands leading with a 31% national average.

EHRs enable highly efficient and remote exchange of health information, remove administrative burdens, enhance productivity, open up new opportunities (e.g., remote health care services), reduce costs, and greatly reduce medical errors. In this manner, they have tremendous value in helping to improve the quality, access and efficiency of health care services. The grand vision of

* Invited submission to the PricewaterhouseCoopers CCE Journal.

EHR infrastructures is the interconnection and reusability of all recorded health information, regardless of where it is stored, so that all relevant health information can electronically flow to wherever it is needed.

Nothing will become of this vision, however, unless critical privacy and security problems are overcome.

2. Privacy and Security

Privacy and security concerns are major impediments to electronic health. If they are not properly addressed, health care seekers will not feel comfortable in participating and health care professionals will face huge liability risks.

Privacy

At the basis of almost all (information) privacy legislation and regulations around the world is the following definition: “Privacy is the claim of individuals to determine for themselves when, how, and to what extent information about them is communicated to others.” (This definition can be made to pertain also to groups of individuals and to institutions, since these may also have privacy expectations of their own.) The definition of Health Canada [2, 3] in the context of personal health information closely mirrors this definition: “Privacy involves the right of individuals to determine when, how and to what extent they share information about themselves and others.” Similarly, the U.S. National Research Council [4] defines privacy in the context of personal health information as “an individual’s right to limit the disclosure of personal information.”

The definition of privacy emphasizes that the control over the release of identifiable personal information should rest with the data subject. Taking away this control from the data subject takes away his privacy. As a consequence, legislation that places this control in the hands of third parties does not do anything for privacy, nor do trusted third party solutions that are unilaterally imposed upon data subjects. Of course, data subjects may volunteer to delegate some or all of their control to others – as long as they have the choice to keep that control to themselves.

The biggest threat to privacy in an EHR infrastructure stems from the secondary use of identifiable health information. Secondary use concerns those cases where information disclosed to one party for a particular purpose is subsequently used for other purposes, without the authorization of the data subject. The opportunities for privacy invasions due to secondary use by insiders are enormous. Organizations with a justified need (according to current widespread regulations) to access health information include government and private health plans, insurance companies, administrators, hospitals, doctors, pharmacies, employers, schools, researchers, data clearinghouses, accreditation and standard-setting organizations, laboratories, pharmaceutical companies, practice management system vendors, and billing agents. Other parties that may seek to obtain patient information include relatives, public health officials, drug marketers, public assistance programs, credit bureaus, and law enforcement agencies. Furthermore, health information custodians may enlist the services of lawyers, accountants, auditors, consultants, volunteers, and a variety of subcontractors.

Indeed, the five threats to privacy in e-health systems that the U.S. National Research Council [4] identified on medical privacy refer primarily to insider attacks. Among the primary threats identified are: insiders who cause accidental disclosures; insiders who abuse their record access privileges; insiders who knowingly access information for spite or for profit; and, vengeful employees. Other studies confirm that the most frequent breaches of patient information confidentiality do not come from unauthorized outsiders, but from uncontrolled secondary usage, accidental disclosures, curiosity, and subordination by insiders.

Studies also reveal that most patients do not trust the administrators of national health services and other insiders in the health care system with the control over their personal health information. In fact, their trust often does not extend beyond their own care providers. Indeed, the unauthorized sharing of sensitive health information may result in a wide range of undesirable outcomes, such as unjustified discrimination by employers, insurance companies, and others. A report by the Health Privacy Project of Georgetown University [5] describes various privacy-protective behaviors currently deployed by patients to protect their privacy: "Patients may see multiple providers to avoid a consolidated record; pay out of pocket for services to which they are entitled reimbursement; and asking a doctor not to write down the health problem or record a less serious or embarrassing condition withhold information; lie; or avoid care altogether." As the authors of the report point out, "Privacy is often portrayed as a barrier to achieving other health care goals. But the opposite is true. People have demonstrated that they need a guarantee of privacy in order to participate fully in their own health care. In the absence of privacy, health care providers often receive incomplete, inaccurate information from their patients, thus compromising the quality of care."

Privacy is also sought by medical practitioners. Many doctors do not like the idea of central parties (such as health insurance organizations) being able to monitor all their actions, since they feel this negatively impacts their autonomy; in many situations, they would prefer to be able to access information on the basis of their role rather than their identity, and they certainly do not want identifiable digital evidence of all their interactions with patients to automatically flow to third parties. Role-based access is also preferred by most researchers, for accessing disease registers and other medical databases.

Ironically, as we will see shortly, the very technologies that are currently being considered to implement important security safeguards may make it impossible for patients and health care service providers alike to escape systemic identification throughout the EHR infrastructure.

Security

Security is generally defined as the extent to which personal information can be stored and transmitted in such a manner that access to the information is limited to authorized parties. Canada's Office of Health and the Information Highway [3] defines security as "The procedures and systems used to restrict access and maintain the integrity of information." Canarie (a not-for-profit corporation dedicated to increasing Canada's economic growth) in its 1997 report [6] on e-health defines security as "The set of safeguards in and around an information system that protect access to the system and the information it contains. […] The purpose of security is to protect the system and the information it contains from unauthorized access and abuse, both from without and from within."

The Health Insurance Portability and Accountability Act (HIPAA) of the United States requires organizations that deal with health care information in electronic form to establish and maintain proper "security safeguards." It requires security mechanisms for dealing with identification, authentication, authorization, access controls, audit trails, accountability, encryption, digital signatures, physical security, disaster recovery, protection of remote access points, protection of electronic communications, software discipline, system assessment for vulnerabilities, and integrity of data.

Note that while security is related to privacy, the two concepts are quite different. Health Canada, HIPAA, and other influential sources on the subject of electronic health information clearly distinguish between security and privacy.

In order to better understand the nature of the relation between security and privacy, it is natural to consider the eight Fair Information Principles codified in 1980 by the Organisation for Economic Co-operation and Development (OECD); these form the basis of most of today's public privacy policy and legislation around the world. The eight OECD principles are: openness;

collection limitation; purpose specification; use limitation; data quality; individual participation; security safeguards; and, accountability. (Several countries actually go significantly beyond this rather minimal set of principles.) Note that "security safeguards" is only one of the eight principles; in other words, security safeguards are necessary to achieve privacy, but they are not sufficient. They do not even contribute to the most urgent pain point in the overall privacy picture, since they deal with protection against access by unauthorized outsiders, while most real-life threats come from secondary use by insiders with authorized access.

Ironically, many of today's commercially available security technologies even have a highly adverse impact on the most important privacy principles, which deal with the ability of data subjects to limit the collection of identifiable information. The deceitful nature of these security technologies is very similar to that of surveillance camera technology: while the risk of being assaulted in public places may go down by placing these areas under permanent camera surveillance, very few people would agree that such an intrusive measure increases their privacy. In fact, many will agree that it creates an electronic panopticon.

3. Privacy-Invasive Security Technologies

When information is in electronic form, the only way to provide adequate security safeguards is through the use of cryptographic techniques. While privacy regulations such as HIPAA do not require specific security technologies, through their strong requirements for security safeguards they have pushed cryptographic security technologies to the forefront.

Public Key Infrastructure

The U.S. National Research Council [4] and other influential organizations around the world have recommended immediate implementation of Public Key Infrastructure (PKI) technology for electronic health applications. It is no wonder, then, that Entrust [7], VeriSign [8, 9], Baltimore Technologies [10] and other leading vendors of PKI technology all have electronic health among their primary target markets. However, PKI technology was never designed for the purpose of access control, which in the context of EHR infrastructures is by far the most important of the security safeguards. As a result, PKI technology is currently being stretched all the way into the realm of access management.

To understand the source of the problem, one must know that the concept of PKI technology was invented in 1978, at the dawn of modern cryptography, for the purpose of facilitating secure message encryption. Consequently, it provides confidentiality of data in transit (through encryption), user authentication (to ensure that a message is encrypted under the right public key), data integrity (to prevent tampering with data in transit), and non-repudiation (proof of the sender's identity). Access control was not on the list of design requirements, simply because it is not of direct relevance in the context of creating a secure infrastructure for sending around encrypted messages. PKI technology can indirectly facilitate access control, however, through its ability to guarantee strong user authentication (by means of digital identity certificates, which bind an individual's public key to his identity). By requiring individuals to provide their digital identity certificates whenever they request access, the service provider can look up any information it wants about the access requestor to make an authorization decision, by using the digital certificate as an authenticated pointer into all sorts of online and offline databases.

There is only so much stretching that can be done, however. Authentication as used for message encryption is a very far cry from controlling the access of authorized parties to identity-related information. Applying PKI technology to the problem of electronic access management in increasingly open environments is like trying to propel a passenger plane using a steam engine; it simply will not fly, and is likely to result in a crash-and-burn.

Indeed, in the overall equation, PKI technology turns out to have devastating consequences for privacy in EHR infrastructures. The fundamental problem is that public keys are unique identifiers: they are the digital equivalent of strongly authenticated SSNs (one cannot deny to anyone the association between one's certificate and one's identity) that are universally unique (each public key consists of several hundred digits) and that travel along with each and every action taken by system participants. Thus, PKI technology forces everyone to transact on the basis of unique identifiers, which are passed on not only to the intended recipient but (inherently to the PKI mechanism) also to third parties that have no business in knowing the details of the transaction. In other words, PKI technology makes it impossible for data subjects and service providers alike to control the flow of personal data, and to limit the opportunity for unauthorized secondary uses of that data; it roots inescapable systemic identification deeply into the infrastructure

Identity certificates that specify a “pseudonym” instead of a real name are not a valid solution: they can be linked and traced as easily and in the same manner as identity certificates, on the basis of their public keys. Indeed, few people would agree that their SSN is a pseudonym. For the same reason, role-based access through PKI (whereby the name field of the certificate specifies a role rather than its holder’s identity) does not allow certificate holders to escape systemic identification.

There is an increasing awareness of the privacy dangers of PKI technology. According to an executive white paper published in January 2001 by the Aberdeen Group, “as currently sold and implemented, PKI is incompatible with the coming privacy era [and] eliminates even the pretense of ensuring user privacy.” Dr. Roger Clarke, an influential IT consultant to the Australian government, in his comments on a major Australian E-Health record project, stated: “The initiative to roll out PKI across the sector was a naive application of ill-conceived and dangerous technology.” The Health Insurance Commission of Australia has already emphasized that they are looking into technologies for “digital proofs of evidence instead of proofs of identity.” As Dr. Peter Swire, then Chief Privacy Counselor to the Clinton Administration, stated in April 2001: “One of the great mysteries is why digital certificates have not spread yet. [...] we need to have both strong authentication and consumer confidence in privacy for electronic transactions. Until that's solved, I'm not sure we're going to see digital certificates spread.”

Privilege Management Infrastructure

Ironically, PKI technology in the context of access management is not even a good security technology: it does nothing to prevent certificate holders from transferring (copies of) their access rights and entitlements to other parties, gives rise to identity theft, and pushes the door wide open to devastating abuses of security holes and errors due to its strong reliance on central databases.

The problems are further compounded by the inability of PKI technology to scale beyond pre-established trust domains: when information is shared amongst different trust domains, the parties that have to make authorization decisions may not be able to obtain all the up-to-date data they need for making an authorization decision. With increasing numbers of individuals and health service providers seeking to share electronic health information, it rapidly becomes infeasible to guarantee the availability, correctness, and timeliness of the data needed to make authorization decisions.

A report [11] by Kaiser Permanente (America's largest not-for-profit health maintenance organization, serving over 8 million members), based on its experience during preliminary design of an enterprise PKI for multiple applications, recommends a shift to privilege management (PMI) in healthcare: “PMI seems a worthwhile approach for managing enterprise healthcare authorization services, particularly those that justify an infrastructure approach because of size or complexity.” The report recommends that “Healthcare organizations should encourage and

become involved in further development of standards in this area. Healthcare organizations should ask PKI and security product vendors about plans for PMI."

With PMI, transactions are conducted on the basis of digital attribute certificates rather than identity certificates. Baltimore Technologies, Spyrus, and other leading vendors of PKI technology have already acknowledged the growing mismatch between real-world needs and PKI based on digital identity certificates, and are actively endorsing attribute certificates; these can specify within them all the attribute data that may be relevant to the service provider. The use of attribute certificates for PMI purposes is even more devastating for privacy than PKI technology, however. On top of the privacy problems inherited from PKI technology, all the attributes within an attribute certificate must be known to the certifying authority and are systematically revealed when showing the certificate.

4. Privacy-Enhancing Technologies

In an interconnected electronic world, there are virtually no grey areas between privacy and inescapable systemic identification. If at the technical level everything is systemically identifiable, privacy legislation becomes virtually meaningless; how can one force participants not to collect identifiable information when they cannot prevent it from being delivered to them?

Privacy-respecting security technologies

The problem can only be solved through the use of security technologies that do not violate other basic privacy principles. In line with the basic definition of information privacy, such security technologies must provide individuals with technical means for controlling themselves the extent to which the personal information they disclose to others can be associated with their unique identity. Until remote mind-control technologies turn from fiction into fact, this implies that privacy-respecting security technologies must, at the very least, allow individuals to de-identify their own personal information before disclosing it.

De-identification of health information is desirable in many situations. For example, anonymity and pseudonymity are entirely appropriate for remote mental health consulting. Also, as we have seen, medical practitioners and researchers have a preference for being able to access electronic resources on the basis of their role rather than their identity. Another example can be found in research on health information, one of the biggest secondary uses of patient health data. Epidemiological research, environmental research, and clinical research can proceed on the basis of de-identified records. Even when linkability with other records of the same data subject is desired, to facilitate linking across different disease registers (e.g., to find correlations), there is no need for identifiability: pseudonymous records will do. In fact, de-identified health data is usually preferable to researchers, because it removes the legal obligation on them to seek consent from the data subject; this avoids wasting time and money to obtain the consents, and prevents a statistical bias in the analysis (data subjects may withhold consent based on the sensitivity of their medical data).

More generally, as the U.S. National Research Council [4] points out: "Only at a few points in the overall health care process is it necessary that the patient's full identity be known. Using (session) identifiers in place of the full patient name […] would preserve patient anonymity more effectively (prevent inappropriate access to patient-identified information) while allowing information to be associated accurately with the proper patient record." In the words of Clarke [12], who notes that patient pseudonyms are already used in such areas as discreet clinics for the treatment of sexually-transmitted diseases and drug-dependency, "Extending these capabilities to additional functions would be of great value to patient privacy. There would be an apparent compromise to the quality of patient care, but this would need to be balanced against the other risks. That choice should rest with the patient."

Controlling who can learn what

It is important to realize that privacy-respecting security technologies are not all about anonymity or pseudonymity towards the parties that data subjects are voluntarily interacting with: they are about the much more general notion of controlling which parties can learn what, as personal information flows through the system from one party to the next. Indeed, individuals typically desire different degrees of privacy towards different parties; their trust in third parties is related to how familiar and comfortable they are with these parties. For example, although a patient may not have a problem with disclosing his personal identity to his doctor when sending a digital attribute certificate, the patient may want to enable his doctor to strip away some of the disclosed data elements (not necessarily his identity) before the doctor passes it on to authorized third parties – PKI technology cannot achieve this.

Since one can always send along additional information, any privacy-preserving security technology offers individuals free choice over the entire privacy spectrum between the two extremes of systemic inescapable identifiability and the maximum degree of privacy attainable through the use of the particular technology. This spectrum is not one-dimensional but multi-dimensional, depending on the number and nature of parties that the individual considers in his privacy considerations.

Two decades of research in cryptography has demonstrated that security and privacy are not trade-offs, but that they are mutually reinforcing when implemented properly. There is no need to rely on an intermediating party that must be fully trusted to not violate the privacy and security interests of the distrustful parties on whose behalf it is intermediating. A fundamental discovery of modern cryptography is that the need to rely on trusted third parties can be eliminated. Specifically, through the use of so-called multi-party computation techniques, any combination of parties can compute any agreed function on their private inputs, in such a manner that the correctness of the result is guaranteed while the secrecy of the inputs is preserved.

In particular, advanced cryptographic building blocks have come out of the cryptographic research community that can be implemented in entirely practical manners to provide privacy and security at the same time. For instance, role-based signing can be securely implemented through the use of privacy-preserving technologies, in such a manner that digital signatures cannot be traced to an identifiable person but only to the role they assumed when signing; at the same time, through the magic of cryptography, in case of a dispute, error, or other mal-event, only those individuals who did not perform the signing can prove their innocence, while the signer cannot repudiate his action. The European Commission [13] stated: “In protection of the privacy of the patient and his right to choose freely his physician or pharmacist and to reveal only the information he decides to, it could be desirable under certain circumstances to have prescriptions, medical reports, diagnoses, entries on the card etc. signed in a way which proves the signature to be legitimate without, however, revealing the identity of the health care professional.” Only privacy-preserving security technologies can achieve this.

5. Ownership Of Health Records

To ensure that patients have control over their health privacy, they should have control over the access to the information. The question is: to what extent?

Personal Health Records

Personal health records have been around for a long time in non-electronic form. Since the early days of medical history, people have carried around their own health records. As physicians started to build long-term relations with patients, and the patients had to visit them instead of the other way around, the control over health records began to shift away from the patient. However, one can still find many examples where patients carry around their own emergency medical

information, typically stored in bracelets or pendants around the neck. For example, expectant mothers, patients in urgent need of organ donation, and military personnel cannot predict in advance at which doctor's office they may end up. Similarly, patients who cannot express themselves (due to mental conditions or otherwise) can benefit from carrying their own health information. (Indeed, the ability to establish trust in a dynamic fashion is also at the heart of the increasing interest in PMI.)

Enter the Electronic Personal Health Record (EPHR). It can include patient identification and contact information, surgical history, demographic information, allergies and other medical conditions, immunization history, family history, organ donor authorization, lab and diagnostic test results, medication and prescription data, health risk indices, provider identification and contact information, treatment plan information, health insurance coverage information, claim status information, appointment information, and so on. In theory, a data subject's EPHR could be the aggregation and unification of all the medical information about him.

With EPHRs, patients can ensure that all the relevant medical information is available at the point of care. They also have greater control over the privacy of their health information, and can ensure that it is accurate and up-to-date. The EPHR can be stored on the individual's home computer, on a portable device (such as a laptop, smart card, or PDA), or on a secure server on the Internet (possibly distributed across multiple trusted parties). With each new consultation with a health professional, the physician is given access to the relevant health information and may update entries in accordance with new findings. In cases where updating is not needed, such as when the patient merely seeks advice on his current condition, the relevant facts can be printed out in hard copy. For health services such as remote consultation over the Internet or participation in online medical chat rooms, relevant information can be sent out in encrypted form.

The idea of giving patients electronic control over their medical information is not new. Gaunt [14] lists over 50 Internet-based EPHR systems. However, these systems are little more than the equivalent of what software like Quicken and Microsoft Money allows consumers to do with their financial data; they help patients track their medical conditions, allow them to review and print information about them, and so on.

More true to the idea of EPHRs are smartcards storing emergency and other health data. According to Waegemann [15]: "During the mid 1980s, […] the vision of patients being in charge of their health information became a leading force. The vision was based on a patient being the connecting entity for all health information. If one would give the patient his health information on a device, he could then bring it with him to providers, thus guaranteeing continuity of care. The practical solution was a patient card […] By the late 1980s this vision failed because of technical card problems, issues with capacities, difficulties with interoperability concerning content and terminology, but most of all due to lack of an infrastructure that allows every provider to record and read cards, even if they were motivated to spend extra time and resources for this exercise."

Controlling who owns what

In spite of the struggles of smartcard-based EPHRs, the shift back towards user-involvement is an undeniable trend. It is driven both by advances in information technology and by health regulations such as HIPAA, which provide individuals with greater control over their own medical information. However, health regulations do not promote in any way the idea of transferring full ownership over health information to its data subject; they typically only provide data subjects with the right to review information and to withhold consent over certain secondary uses.

Indeed, from the perspective of the health care professional, it is one thing for a patient to be able to view his own health information, in order to be informed and to check for errors and out-of-date information; it is quite another thing for the patient to be able to add, delete, modify, or

prevent updating of arbitrary data in the EHR. With medical liabilities being as they are, and most patients not being professionally qualified to make informed modifications to health data, it is no wonder that medical professionals are very reluctant to rely on patient-owned health records; they need to be able to verify the credibility of the source(s) of the health data they rely on.

A report [16] by Cisco Systems states: "Many providers consider the records in their systems to be their property, while patients argue that their medical information is their own. A distinction is often made between ownership of the physical record and the right to access or duplicate data that are stored in it. Policies on health data ownership differ substantially between delivery networks, states and indeed, globally." Larkin [17] notes: "Whether patients or physicians provide information for the record is a subject of often intense debate." Adding to the confusion is the fact that, in contrast to physical objects (where the ability to view, destroy, and modify an object all lie with the person holding it), it is not easy to define the notion of ownership of electronic information.

The big challenge, then, is to find a solution around the complex problem of ownership of health information.

6. SPACER

One possible solution to this problem is outlined here, in the form of an integrated EHR architecture called SPACER (for "Secure & Privacy-enhanced Access Control for E-health Records"). With SPACER, EHRs can be stored on smart cards, PDAs, laptop computers and other portable devices, or on open networks such as the Internet. SPACER facilitates automated sharing and synchronization of certified health information between local and remote health information in accordance with application-specific rule sets.

Security and privacy features

SPACER allows EHRs to be securely managed by both the data subject and his health care professionals, in a manner that simultaneously protects the data subject's privacy interests, the professional's liability interests, and the legitimate interests of researchers and other third parties:

- Each patient can jointly manage his health information together with selected physicians. A record can be managed electronically as one logical entity, even though different parts may reside in different physical locations; those with legitimate access rights might not even realize the dispersed nature of the data they see. Each party with access rights can be assured that the data entries on which it relies have been entered by authorized parties, through either role-based or identity-based digital signatures. In this manner, health care service providers can effectively maintain partial ownership of a data subject's health information; not even those parties who physically control the storage of health information can modify, delete, or add to it without authorization.
- By providing patients with a single tamper-resistant smart card, health care providers can maintain even greater control over their own contributions to a the patient's EHR, since the card can further limit the patient's ability to manipulate entries. SPACER allows for literally billions of digitally certified EHR entries, that may come from different health professionals who do not trust each other, to be securely managed using a single 8-bit smartcard. The storage and computational burden for the smartcard can be off-loaded almost entirely to a patient-controlled device (such as a PC, a laptop, or a PDA), while preserving all the smartcard's security benefits. Card can be issued to patients by a central entity that cannot compromise the legitimate privacy and security interests of patients and health care providers that ride along on its added security; amongst others, different applications using the same smartcard can be fire-walled through the application

software running on the user's trusted computer, rather than having to trust the smart card issuer.

- At the same time, data subjects are able to selectively disclose certified health data to third parties in anonymous or pseudonymous form (with or without certifications). They can also delegate the right to do so to their doctors (e.g., to over-ride protections in emergency situations) or to third parties (e.g., for research purposes) in the form of digital delegation authorizations. These authorizations can be either limited in time or limited in the number of authorized uses. SPACER also provides multiple protocols for gaining access to electronic health information, with varying levels of active participation from the patient and his health service providers.

Through this fine-grained multi-party rights management approach to electronic health information, health professionals can ensure that the data they rely on is of high quality. SPACER in effect creates a continuum between health records maintained by physicians and health records maintained by data subjects, seamlessly unifying the two approaches and covering the entire spectrum of possible rights management settings. In the SPACER approach, the issue of where the health data resides hardly matters anymore; it is all about who has electronic access control to which parts of the record.

SPACER even enables outsourcing of all the security functions to specialized third parties, in a privacy-preserving manner. Specifically, specialized Credential Issuers can digitally certify health information on behalf of health care professionals without being able to learn that data, and Credential Status Providers can validate certificates without being able to learn the identities of patients and health service providers. In this manner, health care service providers can outsource core tasks related to digital authentication and authorization to specialists, without having to provide them with competitive data or patient information for which they could incur legal liabilities.

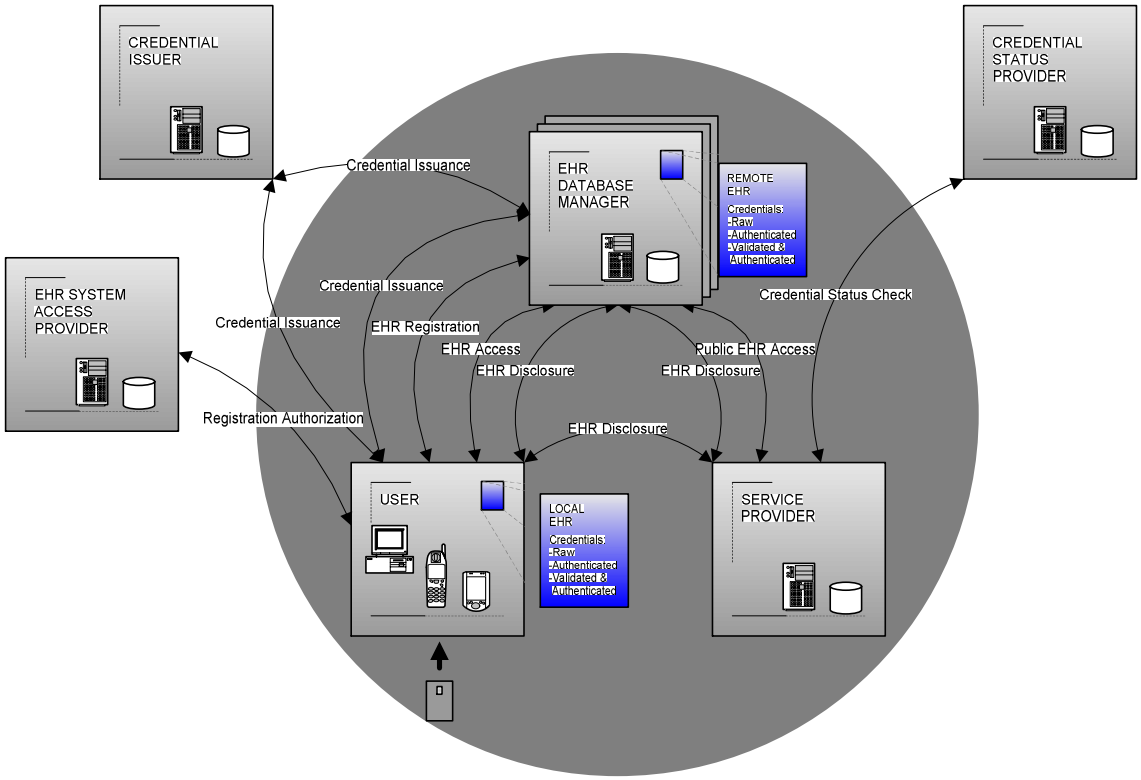

Schematic overview of the SPACER architecture

Even the role of the tamper-resistant smartcard can be outsourced, thereby circumventing the logistical problem of securely distributing tamper-resistant devices to patients and making it impossible for patients to bypass the security of their cards (at the expense of requiring the managed card service provider to be involved in the authorization of all transactions).

Digital Credentials

To achieve these and other properties, SPACER leverages the unique powers of so-called Digital Credentials, a technology designed specifically for the purpose of electronic access management. Digital Credentials are basic cryptographic constructs, much like digital signatures but with much greater functionality. A detailed description of the technology is outside the scope of this paper; for an overview and technical details of the Digital Credentials technology, the reader is referred to an earlier CCE Journal publication [18] and the references it therein.

For the present purpose, it suffices to think of a Digital Credential as a hybrid form of a digital identity certificate and a digital attribute certificate, encompassing both as degenerate cases and offering a multitude of security, privacy, and performance benefits. SPACER makes use of Digital Credentials in three basic ways: to authenticate the data entries of health records, to serve as authenticated pointers to records, and to provide digitally signed audit trails. For a discussions explaining the design considerations behind the SPACER architecture, see Brands [19] and Brands and Légaré [20].

SPACER promotes the notion of standardization of record entries, to drift away from free-form text entries to the extent possible. Ideally, a graphical user interface would allow the selection of entries for insertion into an EHR from a predefined roll-down menu. Not only would this increase efficiency, reduce interpretation errors, and virtually eliminate typing errors, but it would also prevent identifiability of patient-disclosed data by third parties on the basis of the uniqueness of free-form data entries.

7. Final Thought

Ricci, the general manager of IBM's Healthcare Industry, describes [21] a health care future where "Security and privacy issues have been resolved. The question of medical record ownership has totally shifted as consumers have embraced ownership of their own Personal Health Records (PHRs) through secured Web sites. While physicians and hospitals maintain their own medical records, consumers now routinely grant limited access to their complete medical records, and can include alternative and home care." Privacy-respecting security architectures like SPACER are inevitable in making electronic health visions like these become a reality. Security and privacy technologies cannot replace legislation, but in an electronic world legislation has little meaning without them.

8. References

[1] "Toward Electronic Health Records," Office of Health and the Information Highway, Health Canada, January 2001.

[2] "Canada Health Infoway: Paths to Better Health," final report of the Advisory Council on Health Infostructure, Office of Health and the Information Highway, Health Canada, February 1999.

[3] http://www.hc-sc.gc.ca/ohih-bsi/res/defin_e.html, last accessed March 3, 2003.

[4] "For the Record: Protecting Electronic Health Information", Committee on Maintaining Privacy and Security in Health Care Applications of the National Information Infrastructure, National Research Council, 1997.

[5] “Exposed: A Health Privacy Primer for Consumers”, by the Institute for Health Care Research and Policy, Georgetown University, Health Privacy Project, December 1999.

[6] “Ensuring Privacy and Confidentiality on Canada’s Health Iway,” Canarie Inc, December 1997.

[7] “Creating a Common-Sense Health-care Security Strategy,” Entrust, white paper, July 2002.

[8] “Healthcare Authentication Services,” VeriSign, white paper, April 2002.

[9] “Securing Electronic Information in Health Care Organizations,” VeriSign, white paper, May 2000.

[10] “Securing e-Healthcare,” Baltimore Technologies, white paper, April 2001.

[11] “Public Key Infrastructure Concerns in Healthcare Settings,” Dave Barnett, Enterprise Planning and Architecture Group of Kaiser Permanente, published by Kaiser Permanente Medical Care Program, February 2000.

[12] “Research Challenges in Emergent e-Health Technologies,” Roger Clarke, panel session notes, IFIP TC8 Conference on 'Developing a dynamic, integrative, multi-disciplinary research agenda in E-Commerce / E-Business', Salzburg, 22-23 June 2001.

[13] “Final Report of the Eurocards Concerted Action Working Group 2,” Europeean Commission, AIM DG XIII, February 28, 1995.

[14] “Initial evaluation of patient interaction with the Electronic Health Record; South & West Devon Health Community ERDIP Project,” Nick Gaunt, Version 1.0, March 2001.

[15] “Electronic Health Records” C. Peter Waegemann, Health IT Advisory Report, 2002.

[16] “Privacy and personal health records: context, issues and challenges”, Jane Sarasohn-Kahn, Cisco Systems, draft version of January 2001.

[17] “Allowing patients to post their own medical records on the Internet is becoming big business”, Howard Larkin, AMNews, November 8, 1999.

[18] “Privacy In Wireless Environments,” PriceWaterhouseCoopers CCE Journal, issue 4, pages 5 – 11, July 2001.

[19] “Secure Access Management; Trends, Drivers, and Solutions”, Stefan Brands, Elsevier Information Security Technical Report, Chez Ciechanowicz (editor), volume 7, issue 3, September 2002.

[20] “Digital Identity Management based on Digital Credentials,” Stefan Brands and Frédéric Légaré, Proceedings of “Credential-Based Access Control in open, interoperable IT-Systems”, Dortmund, Germany, October 2, 2002.

[21] “Future of healthcare: 2012,” Russell J. Ricci, IBM Healthcare Industry, page 2, June 2002.