

Advanced Database Systems
ECS 289F, Michael Gertz, PhD

Medical Database Security

David B. Hill, MD
November 23, 1998

Introduction

What do people tell their doctor? What do people expect their doctor to tell others? What do people expect their doctor to *not* tell others? This last question is one that is becoming more worrisome in the minds of the public. Needless to say, this concern is driven by the exponential increase in value of information that comes with the ability to rapidly and transparently connect one computer with another via the internet.

Any paper that discusses security must acknowledge that security is a somewhat general term that typically is used to denote three interrelated concepts: confidentiality, integrity and availability. This paper shall restrict its focus to just the first component, but only because it serves as a foundation for all three aspects. The concept of confidentiality in the medical record is a long standing one, to say the least. Indeed, one of the oldest portions of the `medical literature` is the Hippocratic oath, which states in part

*ÓWhatsoever things I see or hear concerning
the life of men, in my attendance on the sick or
even apart therefrom, which ought not to be
noised abroad, I will keep silence thereon,
counting such things to be as sacred as
secrets.Ó*

The Medical Record: What's In It?

But why should people care about maintaining the privacy of their medical record? To answer that question, it is appropriate to look at the contents of such a record in order to explore the ramifications of unfettered public access. Conceptually, the medical record is thought of as a series of dated notes written by physicians on an encounter-by-encounter basis, augmented by reports from laboratories, imaging centers, letters from other consulted physicians. (In addition, the medical record also contains nursing notes which document vital signs, assessments of condition, specific treatment plans for implementing the orders made by physicians, treatments performed, medications given, etc. Despite the acknowledged value and import of nursing notes, this paper will defer further comments on them in the name of clarity and lack of expertise by the author.)

Each of these (physician's) notes is composed of four (4) broad categories of information which model the thought process used by doctors to help organize their thoughts. In order, these sections are subjective information, objective information, assessment and plan. This so-called SOAP format can be used for short simple visits or expanded to a full H&P (History & Physical: the document used when a patient first establishes care with a physician or is admitted to a hospital.) As inferred from the names, the first two sections are where the 'raw' data are stored, the third is the place where the diagnoses are kept, and the last contains the proposed treatment plans (nursing orders).

The history, also known as the subjective, is all of the information that a patient tells the physician, plus information that the physician can obtain from the review of any other source (e.g., old charts, letters from referring doctors, etc.) It has its own internal format consisting of the following:

- Chief Complaint
- History of Present Illness
- Allergies & Medications
- Past Medical History
- Past Surgical History
- Review of Systems
- Family History; Parents, Siblings, Children
- Social History
 - Occupational History (current & past)
 - Sexual History (monogamous, polygamous, straight, gay, bi?)
 - Marriage & Children (number of marriages, outcomes.)
 - Recreational Drug Use (Alcohol, Tobacco & 'Street Drugs')

As you can see, a thorough history involves quite a bit of personal information. Much of this extensive background information has two uses. In addition to providing clues for diagnosing the current issues, the psychosocial aspects of the history are invaluable in determining plans for treatment, since all effective interventions require the understanding and active cooperation of the patient. Knowing as much as possible about the patient's lifestyle and social environment is vital in selecting a treatment

plan that the patient will be willing to take on.

The second part of a medical note is the objective section. Here is where the physician records his/her direct observations about the patient. These observations include everything from a general observation (~Mr. Jones appears disheveled, has not bathed recently and smells of stale beer~) to the results of specific physical examination findings (~Gynecomastia is present. Abdomen is tense and distended, caput medusae are noted as well as numerous spider angiomas.~) In addition, laboratory work (~Blood alcohol level = 0.25 mg%~) and the results of any imaging studies (~CT scan demonstrates cirrhotic liver~) or other specialized procedures are traditionally placed in this section.

Armed with the information listed above, the physician is now able to document one or more conclusions about the current state of the patient and list these as the assessment(s) or diagnoses (~Mr Jones is intoxicated, shows signs of chronic, severe alcohol abuse and has end stage liver disease, presumably secondary to the alcohol.~) Depending on the style of the institution and of the physician, there may or may not be a narrative section here which explicitly documents the thought process used in arriving at the diagnoses.

The final section consists of the plan that the physician (and patient) feels will be most likely to achieve their common goals. Most often these goals are healing and/or cure of disease, but sometimes the goals are less lofty. (~#1 blood serology to rule out concurrent hepatitis infection. #2 Admit to the inpatient ward for detoxification. Will use IV hydration, nutritional supplements, and Valium for seizure prophylaxis. #3 Contact family members re: prognosis. #4 Arrange outpatient detox therapy if patient survives, and is willing to go that route.~)

Example of Data Flow

Now that we have looked at what is in a medical record, let us turn our attention to who can lay legitimate claim to being able to see it. Under today's ~system~ of medical health care delivery, there are a great many people who have this privilege. The following example is taken from[1] .

^Alice~ is a fictional person, but was constructed to be representative of an average middle class American citizen. As portrayed in the reference, she is married to ^Bob,~ is in her middle 20's and has one child. Her health insurance is through Bob's company which is large enough to take on the fiscal risk of being self insured. Alice is relatively healthy, but does have mild anemia (not unusual) and high blood pressure (unusual for a woman of her age.) In addition, she had some complications during her pregnancy which caused her primary MD to refer her to a specialist for assistance in management. Because Alice's high blood pressure is somewhat unusual, she was asked to participate in a drug study.

In order to try to maintain health care costs, Bob's company has directed its employees to seek medical care from a provider who is a member of a particular PPO (Preferred Provider Organization). A PPO is an organization which has agreed to charge lower fees (to Bob's company) in exchange for the increase in patient traffic through their doors. This PPO later finds itself bought up by an HMO (Health Maintenance Organization) which is a larger organization where management places restrictions on the employed physicians in the type of care and medications that they may use in diagnosing and treating patients. Another patient who has been a member of that HMO discovers that she had been denied treatment that Alice had received from the prior PPO, and as a result has brought a lawsuit against the HMO claiming she did not receive proper medical care.

Who Sees / Has Alice's Chart?

Well, now that we know a little about Alice, let's see who else gets to know some of the bits and pieces about her. As daunting as the following list appears, keep in mind that no one person or organization actually has access to *all* of the information that pertains to Alice.

Alice's *primary physician* obviously is one of the key individuals who has access to the vast majority of her clinical and personal medical information. In addition, there are the *clinical laboratories* which perform the lab studies needed by the physicians. Although they do not see any of the history and physical data, some conclusions can be drawn from the tests ordered, in addition to the data inherent in the results of those tests. Similarly, the *local pharmacies* can deduce a broad understanding of a patient merely by examining the types of drugs and the usage patterns (dosage,

frequency). Again, they are bereft of the details from the history and physical exam, but since they have Alice's identity (something the labs don't necessarily have) they now have some insight into a specific person's medical condition.

When we created Alice, we gave her a condition requiring a *consulting physician*. This individual will typically need to collect his/her own history from Alice, but if the consultant had access to the same health care information system as the primary physician, then presumably everything that was recorded by the primary doctor would be available for the specialist. If the specialist was in a different health care system, then his/her clinic or *other local hospitals* would now have a different chart on Alice and begin compiling their own set of details on her current medical condition.

So far, the people who have access to Alice's data have been health care providers or allied professionals. Let's look at some of the other people with access. In order to maintain its credentials, hospitals undergo periodic review. During these reviews, patient charts are pulled at random and examined to determine if the physicians and nurses have provided appropriate and adequate levels of care to their patients. In these audits, the *auditing team* has full and complete access to any and all data. We mentioned that Alice had a baby in her scenario, thus the local *state bureau of vital statistics* shall also be privy to at least a portion of Alice's record.

One of the major weaknesses in the case for patient privacy comes from the fact that in this country, health care is almost always paid for via some other '3rd party payer.' This 3rd party is often an insurance company, or a state/federal program for indigent patients. In either event, there is a strong financial interest on the part of the payer to verify that they are paying for services which were needed. In other words, they tend to ask for details from the physician's notes to 'prove' that the interventions performed by the physician were medically necessary, not to mention that they were covered by the contract of what would and would not be paid for. (Patient expectations for privacy in this regard are typically dealt with by way of standard legal language that the patient must sign before s/he ever gets to see a physician.) In Alice's case, who are these 3rd parties? Of small concern is the company that underwrites her medications. Of larger concern

is the fact that the company that underwrites her clinical and hospitalized care is her husband's employer! Immediately, it is obvious that there is the potential for a significant conflict of interest. If, for example, a personnel officer notes that Alice has developed a chronic (read: ^expensive~) condition, then it stands to reason that one way the company could save money would be to put pressure on Bob to leave the company. The situation only gets worse if Alice had insurance provided by her own company instead of through her spouse's.

Insurance companies do not operate in a vacuum. There exists an organization called the *medical information bureau* which exists in order to record what conditions various insured individual have so that when they change insurance companies the new company can determine if there are any reasons that coverage should be denied or covered at a higher rate.

Recall that Alice was being followed as part of a long term medication study. This implies that her records are available to the researchers who are conducting the study. (There exists a trend among researchers to examine only statistically abstracted data, without identifying information. Although it is possible to use even this kind of data to determine facts about a specific person[2] , it is a step in the right direction.) In addition, Alice's records have been subpoenaed as part of a lawsuit against her new HMO. Perhaps one of the most concerning aspects of this wrinkle is (as I understand it) if her chart is entered into evidence in the trial, then the chart becomes public domain information as part of the legal proceedings.

Finally, let us take note of how much of this information is moved from one place to another. Most medical records are not yet in electronic form, but almost all of the financial aspects of medical care have been electronic for years, thus large quantities of data are moved via tapes, modem connections, and internet connections. Unfortunately, the use of encryption within the medical community is practically non-existent.

User Profile

Before proceeding to a discussion on the databases themselves, we should also look briefly at one of the primary users of the medical database, the physician. Although certainly not true in all cases, it is helpful to think

of physicians & nurses as being ^computer-phobes.~ If you assume that they can't type, can't recall non-trivial passwords, refuse to go to training classes, refuse to use systems with more than 5 seconds of latency, and in general simply hate computers, then you have a good initial understanding of the kind of user your system must deal with.

Another aspect of the physician-user that tends to get overlooked is the assumption that s/he is relatively static within the organization. Certainly, this may be the case in some practice conditions, but it is not always true. For example, internal medicine teams within my teaching hospital routinely admit 14 patient per call cycle (every 4 or 5 days). With a steady state load, this many patients must also be discharged during the same time frame. Each team has a minimum of 4 MD's, plus 1 to 3 students, all of whom change roles and go to different teams and/or hospitals approximately once per month. Clearly, keeping track of which doctor is seeing which patient would represent a significant administrative cost, not to mention delay, if done manually.

Another example demonstrates even higher `patient flux.` Patients who need to be seen in an urgent care setting are typically evaluated and treated by one physician (who ever is on-call for urgent care that shift) and followed by another, (patient's primary MD, or who ever is on-call when follow-up is needed.) In this case, the patient/physician pairing is not made until the time of care. Further complicating this scenario is the question of `who should late results be given to?` If the system tracks outstanding labs, etc., by the identity of the ordering doctor, is that necessarily the best person to be given results that come back the next day? Often, the answer is `no.`

System Profiles

Having looked at the typical user, we should now examine the typical systems that they use. A typical early type of electronic medical record should not really be called a `system.` In many institutions, the electronic patient record is derived as a side effect of having computerized billing systems. These systems were initially deployed to track the complex negotiations of patients, providers and payers. One of the major strengths

of these billing systems is that individual patients were strongly identified, i.e., these systems used large sets of demographic data plus the patient's name in order to attempt to uniquely identify each patient. Having strongly identified the individual patient, the system then attached an abstracted disease code (e.g., ICD-9-CM^{1*}) as part of the billing process. With patients and (rudimentary) disease lists linked together, these systems were pressed into service to begin giving information to the health care providers. Since these collections of software were not really intended to provide health data, I refer to them as 0th generation systems.

Extending this nomenclature, systems actually designed to provide medically relevant patient data should be referred to as 1st generation systems. The most famous of these began at Massachusetts General Hospital and was written in a somewhat obscure language, MUMPS, that eventually became an ANSI standard, M. Systems written in MUMPS are inherently hierarchical in nature with a patient-centric point of view. Initially text only, they have recently begun branching out to provide additional forms of information. Being hierarchical, they suffer from the inability to change their point of view. This lack of flexibility makes them difficult to use to perform research, as it very difficult to search the patient database for sets of patients based on disease, treatment, outcome, etc. Nonetheless, these systems are still tremendously valuable at doing what they were designed to do, namely to provide physicians with patient information in order to treat that patient.

The 2nd generation of systems are those which are based on a relational database, either an off-the-shelf product such as Sybase or Oracle, or a custom engine designed from scratch for use in a medical environment. These products have suddenly exploded onto the market place over the last 2 to 4 years. In the vast majority of the sales literature being used, these system still show their roots in patient billing, now known as ^cost capture.~ Administrators have determined that if you have the actual care provider enter his/her findings, assessments and plans directly into a computerized system, then that system has all of the required information to produce direct billing for every aspect of care that has been

^{1*} International Classification of Disease, 9th edition, with Clinical Modifiers. (1977)

rendered. This represents an improvement in business efficiency as it eliminates the need for clerks to sift through patient records looking for information with which to generate and/or support billing. In addition to improving cost-capture, these systems should (but don't always) reduce some of the repetitive data entry that typify current manual (paper based) systems. It is of academic interest to note that none of these systems claim to actually benefit the patient population as determined by a formal randomized, clinical trial. Relational DB systems are being marketed to small clinics staffed by a single MD, to full tertiary care teaching hospitals with 500+ beds, and hundreds of providers.

Finally, the 3rd generation systems are those which are using HTML or XTML interfaces to commercial browsers, such as Netscape Navigator. These may be based on either hierarchical or relational backends, and are simply taking advantage of the huge base of previously installed generic computers with browsers as a "free" front end. The concern that these systems raise is the possibility of wide spread loss of confidentiality. Although security is generally not a selling point in electronic record systems (see below) the marketing of web enabled front end based systems is a refreshing exception to that trend. These systems explicitly acknowledge the tremendous damage that could be wrought if patient data were made readily available on the internet[3].

Despite the appreciation for patient confidentiality concerns expressed above, several vendor representatives have expressed (through personal comments) that security features are highly sought after in the sales process, but are uniformly ignored when the software is delivered / installed. These vendors rapidly reached a consensus that when such software is purchased by a hospital or clinic, there has been a great deal of high expectation generated regarding the new software. Simultaneously, the Information Services department is trying to bring the new system on-line with as little delay, disruption, and chaos as possible. Meanwhile, users who are used to the "old system" may be quite skeptical of having to change the way they do their work, and will use any perceived fault of the new system to resist using it further. These pressures quickly lead IS to disable any possible security feature that could delay / interfere with the users' acceptance of the new system. Sadly, once security features are taken down,

they can only be re-instated with great difficulty due to extreme user discontent. The bottom line is that security is the last thing that users are interested in.

Secure Databases

Having explored some of the peripheral issues concerning the contents and users of a medical database, we can now examine the idea of using a secure DBMS. For the purposes of this paper, I use the term "secure DBMS" in the same sense used by Landwehr[4] : a DBMS which contains the data of interest, plus metadata regarding the sensitivity of individual data elements, metadata on the security classification (level) of authorized users and only permits users to retrieve data which their classification allows.

Work in this direction has been explored by Pangalos, et al, at the University of Thessaloniki in Greece[5] , [6] , [7] . Pangalos emulates a secure DBMS by using Oracle version 7, and explicitly providing additional attributes for each portion of target data to represent its sensitivity. Separate tables contain users and their security classification, and specific views are used to implement additional selection restrictions on queries in a manner not visible to the end user. This system has been fielded in an actual working environment, namely the 700+ bed University Hospital in Thessaloniki. At run time, queries from users are modified to add additional restrictions on retrieved data based on physical location of the terminal, the application used to make the request, and the identity (security classification) of the user.

Pangalos avoids the problem of cascading authorization inherent in the discretionary access model by not permitting physician users to grant access to other physicians or nurses. He defines three groups of users; technical users (DBMS administrators), hospital administrators and medical staff. The technical users have direct access to the DB. Medical staff are restricted as described above, but administrative staff do not have the location of the request factored into their queries.

Despite the fact that this project has actually been fielded and is in a working environment, there exist several problems which need to be addressed. The first of these is the method of access control. The use of a separate table delineating users and their classification, plus the inability of

those (medical) users to grant or revoke patient access to other users indicate that Pangalos uses a mandatory access control (MAC) policy. As mentioned above, the high degree of 'patient flux' implies a high administrative cost with such a system. Furthermore, the nature of medical care is better modeled by the concept of medical roles vs. identity of the physician providing care (recall the explanation of the urgent care clinic.) Pangalos specifically acknowledges the role based model of access control (DAC) as being more appropriate for medical environments and lists that as an area of continued research.

A second, and to my mind more concerning, issue pertains to the ability of a multilevel DBMS to provide 'possible explanations (cover stories) for unavoidable observable information that would otherwise lead to partial or complete inference of sensitive information.' In other words, a multi-level DBMS can provide different answers to the same query, based on who is asking the question. Although admirable from the security point of view, this solution to the inference problem (discussed below) is not acceptable for several reasons.

Firstly, the technique of spreading misinformation is a military tactic which owes much of its success to the subsequent use of the incorrect information. Errors in the administration of such a system could easily result in (authorized) users being provided 'cover story' information instead of the true value contained within the DBMS. If this misinformation were to be used to make decisions regarding patient care, it could result in increased morbidity and/or mortality ' exactly what military tacticians seek to accomplish and what health care providers wish to avoid. Second, confidence in such a system would be sharply eroded since users would be aware of the fact that the system had the designed capability to provide disinformation. This could result in the *de facto* boycott of the system since it is preferable (to some people at least) to work in the confidence of ignorance rather than the shadow of uncertainty.

Lastly, the technique will not solve the problem of inference anyway. This is because the system administrator (the person who creates the misinformation) has no way of knowing which facet of the medical record each patient would wish to protect. In addition, the system administrator *can not* make a coherent set of misinformation that will survive scrutiny by a

determined opponent. This last issue is part of what is referred to as "the inference problem."

The Inference Problem

This problem of inference is particularly tricky in the context of medical information. One of the hallmarks of biological systems is their extreme degree of inter-relatedness. This complexity makes the concept of normalization of the database impossible to accomplish in any meaningful way from the point of view of a knowledgeable user. This is not to say that the data can not be normalized in a technical sense - it is certainly possible to reduce data so that there is no redundancy in the stored tables, but it is not possible to remove the redundancy imposed by the physiology of the patient which that data represent.

Given this degree of inter-relatedness, it is certainly possible for a user with a low classification to obtain information which will permit him/her to infer the value of a datum at a higher classification. For example, it has been shown that information regarding a single patient can be obtained by properly phrased queries involving only abstracted statistical data, and that this data leak can not be closed by the use of authorization restrictions[2].

Referring back to the section discussing misinformation in a medical DBMS, it is not possible for such a system to provide a convincing "cover story" by changing a significant piece of medical data without altering all of the other data which rely upon the index datum or are otherwise influenced by it. (e.g., you can't simply change the diagnosis from "kidney failure" to "liver failure" without altering the history of the presenting illness, the physical exam findings, the laboratory data, any images which show the kidneys and/or liver and the therapeutic procedures for treatment.)

The inference problem becomes even more intractable when you consider multiple inter-related databases, such as those in a hospital, an out-patient clinic, a pharmacy and insurance clearance house. Assuming that each of these DBMS have sufficient demographic information to identify an individual patient, it is possible to determine classified facts in one DB by relying on non-classified information contained in another. For example, if the hospital has classified that patient Jones is HIV(+), and

further suppresses that information by blacking out individual test results that are consistent with HIV, it would still be possible to deduce that Jones was infected by simply inspecting the medications s/he is taking since there is one and only one reason for taking drugs from the class which targets the HIV protease. Even if these drugs were blacked out from the pharmaceutical DB, it is still potentially possible to deduce Jones' infection if s/he has been placed on a fairly common drug protocol designed to prevent infection from *Pneumocystis carinii*, an organism which is harmless to people who do not have AIDS, but in AIDS patients is treated with the same antibiotic used for ear, sinus and bladder infections. (The difference is the dosage and time of use in the case of HIV infection vs. the other cases mentioned.)

This problem was explored in a highly technical paper by Kinoshita[8]. Although difficult to follow, the conclusion reached is that if multiple multi-level databases are used, it is not possible to prevent the inference of a classified datum without raising the classification of all inter-related data in all of the DBMS's. In other words, you can't rely on omission or misinformation to prevent data leakage by inference.

Filtering Requests

One avenue for partial solution of the inference problem is suggested by Wiederhold, et al, in [9, 10]. They propose a method in which queries from external, but related, DBMS's are checked against a set of rules before being executed. These rules would be defined by the local information security officer. Any queries which were denied would, of course, be logged. Queries which could not be resolved by the rule base would be flagged for human attention. The authors propose that these exceptions would then provide the security officer with information (and incentive!) to further refine the rule base to reduce the number of exceptions so generated. Unfortunately, the authors provide no formal description of how such rules could be constructed.

Although the focus of these papers is on queries generated by other DBMS's (e.g., the pharmacy computer querying the hospital DB) it seems a reasonable course of research would be to investigate if AI techniques could be used to spot 'data mining' queries or those used to make inferences

about classified data. For example, individual users working through an application can only operate at human speeds. An application which is requesting more than 1 record per second could be assumed to be some sort of software agent masquerading as the logged in user, since no human can generate requests at that rate. Another "AI" rule could be to look for multiple statistical queries from the same logical user, where the queries have only few differences from previous queries.

Again, these techniques would require a great deal of research before it could be determined if they would provide any sort of meaningful security.

Auditing As A Method of Deterrence

The text *For The Record* [1] covers a great deal of material with respect to all aspects of medical record security. One concept that was well covered in the text is that most of the improper access to medical records is by users who have legitimate access to the system, but are using it in an inappropriate way. In the current literature, this concept seems to be largely overlooked, with the bulk of security work being aimed at providing a secure perimeter against external threats, but largely ignoring internal personnel. As an example, Yeager and Kindle[11] document in great detail the security arrangements made at a military clinic. Their entire article, with the exception of one paragraph, is concerned with protecting their data against external threats.

We have looked at some methods of defeating improper access by otherwise authorized users, but so far, none provide a simple nor completely effective method of effecting confidentiality. *For The Record* advocates a somewhat novel approach that combines the good manners your mother taught you ("don't snoop") plus a popular American weapon ("or I'll sue you.") The concept is simple; in addition to granting all patients the right to see their medical record (much like they have the right to view their credit record) patients should also be granted the right to see *who has viewed* their medical record. In addition, health care organizations need to produce clear, unambiguous policies that state that any employee is free to access any health record as needed in the line of duty, but only in the line of duty. Furthermore, any and all accesses to any medical record would be

recorded and subject to audit (both internally and by the patient / patient's agent.) Lastly, a clear policy which outlines the penalties associated with improper access must be made clear to all employees. With these ingredients, the temptation to browse the charts is tempered with the knowledge that such browsing is very expensive.

As an anecdotal story (provided at the Computerized Patient Record Institute meeting 11-6-98) describes a case in which a local celebrity was shot and transported to a community hospital in Florida, where he later died. During his last few hours within the hospital, several employees made use of the medical record system to call up his case, simply for their own curiosity since the victim was a "celebrity." They did this despite explicit hospital policy to the contrary. After the case was closed from a medical point of view, the (4?) employees were investigated, and eventually dismissed for violation of hospital policy. As a student of human behavior, I am willing to bet that similar infractions of patient confidentiality will be quite rare at that hospital.

Like the other technologies described above, simple capture of audit trails will not be sufficient to make significant changes in human behavior. For example, the hospital I work in routinely captures the ID of persons accessing transcriptions of dictated reports, as well as accesses to laboratory information, but no person ever reviews those logs, thus they serve no purpose as a deterrent.

Conclusion

In this paper, we have examined the contents and users of medical databases. Although slow to reach the medical community at large, electronic medical records are rapidly making progress in becoming the standard of care for patients in this country. Security techniques which have served the military infrastructure are being used as initial points of development for the security (confidentiality) of patient information, but there significant differences in the way that military data and medical data are used, thus military techniques do not adapt well to the new environment.

In order to provide a meaningful level of data protection, further research is needed in the areas of role-based authentication and the use of

improved audit and analysis tools to help human security officers make reliable and meaningful evaluations of both internal and external threats.

One avenue which is currently available is the use of access audits coupled with `after the fact` legal sanctions against violators. These techniques should provide not only a reduction in the number of inappropriate queries against confidential data, but also the means to punish those who disregard the privacy of the patient.

References

1. CSTB, *For The Record: Protecting Electronic Health Care Information*. 1997: National Academy Press. 264.
2. Miller, M. and J. Cooper, *Security Considerations for Present and Future Medical Databases*. *Int'l J of Bio-Medical Computing*, 1995. **41**: p. 39-46.
3. Halamka, J. and C. Safran. *Virtual Consolidation of Boston's Beth Israel and New England Deaconess Hospitals via the World Wide Web*. in *AMIA*. 1997. Nashville, TN.
4. Landwehr, C., *Formal Models for Computer Security*. *Computing Surveys*, 1981. **13**(3): p. 247-277.
5. Pangalos, G., *et al. Development of Secure Medical Database Systems*. in *DEXA*. 1994.
6. Pangalos, G.J., *Design and Implementation of Secure Medical Database Systems*. *Medical Informatics*, 1996. **20**(3): p. 265-277.
7. Pangalos, G.J., *Secure Medical Databases: Design and Operation*. *Int'l J of Bio-Medical Computing*, 1996. **43**: p. 53-60.
8. Kinoshita, H. and S. Tsujii, *Information Security of Database Networks*. *Systems & Computers in Japan*, 1990. **21**(13): p. 22-29.
9. Wiederhold, G., *et al. A Security Mediator for Health Care Information*. in *AMIA*. 1996.
10. Qian, X., *et al., Trusted Interoperation of Healthcare Information (TIHI)*, . 1996: NSF Challenge Grant PI Meeting. p. 2.
11. Yeager, R. and M. Kindle, *Information Security at the 15th Medical Group*. *J of Healthcare Information and Management Systems Society*, 1998. **12**(1): p. 39-50.