December 2001

Old Tricks for New Dogs: Some Methods for Improving Internet Agent Security

Steven Walczak
University of Colorado at Denver

OLD TRICKS FOR NEW DOGS: SOME METHODS FOR IMPROVING INTERNET AGENT SECURITY

Steven Walczak
University of Colorado at Denver
salczak@carbon.cudenver.edu

Abstract

Electronic information retrieval is becoming a necessity for business. Agent technology promises to provide a facile means for intelligently searching the internet. However, any agent or other electronic information placed onto the internet poses a security risk, especially in information sensitive domains such as medicine and law. Encryption and firewalls are traditional methods for increasing internet information security, but are these techniques sufficient to adequately protect electronic information carried by internet agents? Need to know and partial information strategies are discussed as methods for improving information security for internet agents.

Keywords: Internet, world wide web (WWW), security, medical information, internet agents

Introduction – Need for Medical Information Retrieval Agents

The world of medicine is racing to embrace new information technology and in particular the e-commerce benefits provided by the internet and world-wide-web (WWW) (e.g., WebMD, http://www.webmd.com). However, exchange of medical information over the internet enables compromise of patient confidentiality. The vast quantities of information available on the WWW and the speed with which new information can be posted and dispersed make the WWW a tool of growing importance in delivery of high quality patient care (Detmer and Shortliffe, 1997; Spath, 2000). Due to economic and performance pressures, medical offices and institutions require user-friendly rapid access to information and electronic medical records is the start of providing such access (Rindfleisch, 1997), but poses corresponding security risks. Physicians (Bazzoli, 2000) and consumers both are searching the WWW for diagnosis and treatment information. Using electronic medical records in combination with intelligent WWW information retrieval will be a critical success factor for medical practitioners.

When using search engines, the aggregation of terms used to guide the WWW search process poses a potential breach of confidentiality risk in that the aggregation of terms may be collected and later associated with a specific patient. An alternative to traditional WWW search engines is the use of agents. Agents are a new technology, enabled by the WWW and intranets, that facilitate searching the WWW for medical (as well as other types of) information. An agent is a software program that autonomously acts on a user's behalf to satisfy a specific task, which in this case is the location or tracking of medical information. Agents may also be used to track user movements on the WWW and are frequently used by e-commerce sites, such as Amazon.com to record a user's activity while at the site.

This article will examine information security measures that will provide relatively safe utilization of agent technology for searching the WWW to locate medical diagnosis and treatment information. The next section, *Background*, briefly summarizes internet communication security risks and agents. The *Old Tricks for New Internet Dogs* section proposes an agent framework for minimizing electronic patient information risks and develops a heuristic measure for evaluating the security risk of compromising confidentiality while using internet agents to perform information retrieval.

Background – Agents and Internet Security

The WWW resides on the internet, which may be viewed as a very large network of interconnected computers. Because of the high interconnectivity between the various computer networks that comprise the internet and similar interconnectivity within each network, placing any type of data onto the internet essentially broadcasts that information to every computer connected to the

internet. Any information placed on the WWW should be considered to be global knowledge within twenty-four hours (Zwart and Resnick, 2000). Common security and data transfer protocols disable reading messages from IP nodes that are not on the intended recipient list, however these protocols may easily be thwarted, enabling "eavesdropping" on the internet and WWW. Internet nodes may be set up to perform "sniffing" or reading of information that was not intended for that internet address (Atkins et al., 1997).

The traditional means for securing electronic data transmitted over non-secure telecommunications lines, such as the internet, is encryption (Hsiao et al., 1979). Transmitted data is transposed to alter its look such that intercepted messages cannot be read easily, at least not without knowing the key. Existing methods of secure socket layer (SSL) technology with 132-bit encryption provide a modest amount of telecommunication security. Use of encryption does not mean that an intercepting party cannot read a message, only that the interceptor will have to spend a greater amount of effort, in decoding the message, to read the contents of the message (Hsiao et al., 1979).

Firewalls act as a barrier between a physician's host computer and the internet network of computers (Atkins et al., 1997; Garfinkel and Spafford, 1996). Other traditional physical security and access security measures should be enacted in the physician's office or hospital setting to ensure that patient privacy requirements are met with respect to local electronic patient records. Examples of physical and access security include: locking the network server's room to prevent entry by non-authorized personnel and password protection on network/computer accounts that have access privileges to patient information.

Agents provide a means for automating search and other tasks for users. Standard encryption methods may be used within agents to provide a minimum level of information security, but again encryption does not mean that the agent's information cannot be read. Therefore additional measures to improve the confidentiality of electronic patient information is needed and discussed in the next section.

For security purposes, only those agents that are categorized as "mobile", able to move around a network or the internet, are of concern (Maes, 1994). Stationary agents that automate office reminder systems (Silverman et al., 1998) or other local tasks will be protected via firewalls, physical security, and access security measures (Hsiao et al., 1979) that should be provided by the medical office or hospital. In addition to mobility, other classifications of agents correspond to their cooperation (multi-agent versus single agent systems) and intelligence (ability to adapt to unforeseen circumstances) (Maes, 1994). Another important aspect of agents is that they may be viewed as an encapsulating control mechanism that enables plug-and-play adoption of new security and other methodologies, such as natural language parsing, as they become available (Bigus and Bigus, 1998; Silverman, 1998).

While cooperation among agents that have been produced by multiple independent systems is a theoretical possibility (Finin et al., 1997), the reality is that agents typically only need to cooperate with other agents that have been generated from the same agent-based system. Should inter-system agent cooperation become necessary, then new identification and verification security measures will become necessary.

Agents may also be programmed to act like an internet "sniffer" (Atkins et al., 1997) and read the links followed by users and the user choices made while on specific web sites (Ip et al., 2000). E-commerce sites, like Amazon.com, typically use site-limited agents to track their customers' preferences and provide personalized service. Passive security attacks enabled by agent "sniffers" are difficult to detect and more difficult to defend against (Hsiao et al., 1979). As long as there is a high enough volume of internet traffic that is directed at a large variety of web sites, agent sniffers will have difficulty in compromising an individual's privacy.

In medical domains, agent technology has seen a surge of application development in the past several years. Agent reminder systems that operate off locally stored protocol information have been implemented in medical domains as stand alone systems (Silverman, 1998; Silverman et al., 1998). Additionally, other agent-based systems have been developed in medical domains to provide user-friendly search engine interfaces (Baujard et al., 1998), to facilitate the transmittal and encoding of multimedia data (specifically laboratory tests) (Della Mea et al., 1999), and to provide expert-like decision support from a local knowledge base (similar to the reminder agents of Silverman) (Della Mea et al., 1999). Most of these agents operate as stationary agents and as such will not require the electronic transmittal security measures discussed in the next section. However, a framework for using multiple mobile agents to dynamically search web sites for medical information that overcomes the problems associated with using existing index-based search engines has recently been introduced (Walczak, 2000).

Searching the internet for information and retrieving related documents is a task that may be automated via internet agent programs and requires that these agents be mobile (able to traverse the internet). When a medical professional requires or desires the latest information from the internet, two variables affect that professional's search productivity. These two variables are knowledge level of the document to be retrieved and familiarity with the web site where the information may be obtained (including mirror sites). If a specific document at a well known web site is desired, then a direct ftp download or web-based

lookup of the document may be performed. If either variable, site or document knowledge, is not completely specified, then some type of search must take place (Walczak, 2000). Typically a physician would be lacking document knowledge for newly created documents available through the WWW or internet that cover diagnosis and treatment of existing patients. Existing search engines suffer from many faults including: inability to provide the latest information, providing too much information (information overload), providing out-of-date or dead links, and providing non-relevant information. Internet agents, by searching the current WWW information space at a targeted list of web sites overcome most of these drawbacks of current web search engines.

Old Security Tricks for New Internet Dogs (Agents)

The use of multiple mobile agents for identifying and retrieving web-based medical information (Walczak, 2000), while satisfying the growing need of physician's for facile interaction with the WWW, presents patient privacy/information security problems. First, all medical information that is carried by any mobile agent is encrypted using standard WWW encryption methods. Although encrypted, recall that encryption does not guarantee security but rather merely makes the agent's information more difficult to read, additional security measures are needed to improve patient confidentiality.

The first security practice for multiple mobile agent implementation is the old adage: "need to know". A problem with using multiple agents is that at some point in time, the information acquired by various agents must be integrated into a single source for the physician to utilize in treating a patient. Typically, this means that each agent will carry not only it's search criteria, but also patient information to facilitate the integration of any returned information. However, searching WWW sites for medical knowledge does not normally require patient demographic or identification information. Therefore, medical agents should only carry medical search criteria and not patient identification information.

The isolation of patient identification information can easily be accomplished by performing all of the mobile search agent's information integration locally. This will require that an integration agent be implemented locally that keeps track of all existing mobile agents and how their information is related to specific patients. The non-mobile integration agent must be subject to and protected by the same physical and access control security measures that are already in place for other locally held electronic patient information (e.g., firewalls and passwords), as shown in Figure 1. Using the local integration and mobile search, sniffing of an agent's information will not lead directly to a loss of patient privacy, since patient identification is performed solely on the local computer system of the physician's practice or hospital.

A side benefit from removing patient identification information from the mobile agents is the resultant reduction in size of the mobile agents. The smaller sized agents will reduce the overall volume of internet traffic that would be otherwise produced by the larger patient specific mobile agents. The drawback is that the local integration agent must now maintain a database of all existing mobile agents and their corresponding patients, which might impact the performance of the local computer system.

Even though the specified mobile agent security protocol of not carrying patient identification information provides improved security, if a sufficient quantity of information via the search criteria is carried it may still enable the identification of the corresponding patient. For example, a physician may treat several hundred diabetic patients and similar numbers of patients suffering from either glaucoma or neuropathy or kidney failure. However, the combination of all four terms into a single agent will narrow the prospective list of patients down from hundreds to only a dozen or possibly even uniquely identify a patient from the collection of medical conditions. Therefore, an additional security measure is required that will prevent indirect patient identification.

The second security method is similar in concept to the "need to know" protocol and concerns the use of partial information strategies. The use of partial information comes from a technique used to enable two people who have never met to identify each other. Each individual is given one half of a uniquely cut object that will only match the other half of the same object, like a jigsaw puzzle piece. Partial information strategies require collaboration amongst multiple agents, similar to collaborative problem solving (Sayeed et al., 1999). The mobile agents should each be given only a single atomic term to search for on the WWW instead of a large collection of terms. Since the individual terms, as used in the example given above (e.g., diabetic, glaucoma, etc.), do not provide enough information to uniquely identify a single patient or small group of patients, the risk of a patient privacy breach is minimized. The integration agent will then assemble the multiple mobile agent information retrievals into a meaningful composite whole.

The partial information strategy necessitates the implementation of multiple mobile agents, each of which carries a single search term, to satisfy the medical information retrieval needs of physicians for specific patients. An integration agent capable of determining patient agent relationships is already specified from the "need to know" protocol and now must be expanded to simultaneously integrate multiple information sources and remove duplicate retrievals (to reduce information redundancy). Fortunately, there already exist many algorithms for determining relevancy of documents when multiple documents are retrieved

via a search (Menczer and Monge, 1999; Tu and Hsiang, 2000) and agents may easily implement new information ordering and integration techniques as they become available (Bigus and Bigus, 1998; Silverman, 1998).

One further mobile agent strategy may be implemented to further prevent adverse "sniffer" agents from attempting to accumulate a sufficient quantity of information to compromise patient confidentiality. As mentioned before, passive attacks like sniffing are very difficult to detect and only large amounts of non-organized web site visits can provide a limited obfuscation defense. Thus, while each of the mobile agents carries only a single search term and does not carry any patient identification information, the pattern of web sites visited by an agent may provide indirect evidence of the agent's search mission. Since the local integration agent must already integrate multiple agents, corresponding to the different search terms for a patient, increasing the number of agents and limiting the sites visited by each agent, thus eliminating any pattern, may provide further information security. A separate agent should be created by the multiple agent information retrieval system for each medical information web site to be visited.

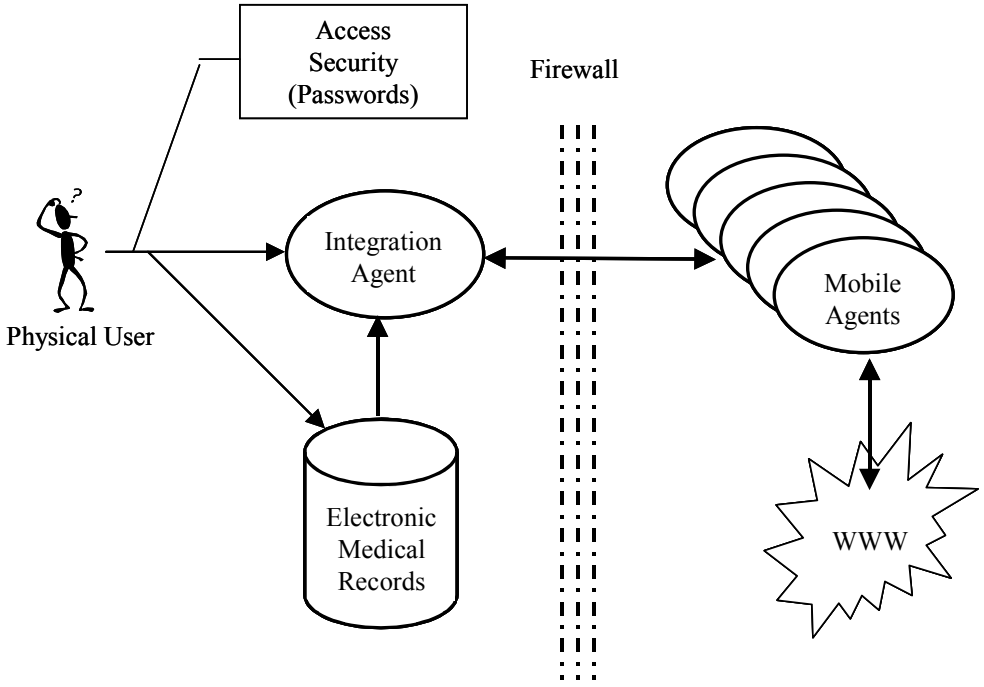

Figure 1. Local and Mobile Agent Architecture with Local Security Controls

Physician's may be queried to provide a list of high quality and relevant WWW addresses or a default list may generated automatically by performing random medical term queries on existing WWW search engines. Each medical information search then produces *n***m* mobile agents, where *n* is the number of WWW sites to be searched by an agent and *m* is the number of medical terms comprising the information retrieval search. The *m* medical terms are normally comprised of noun or verb phrases that are specific to the signs, symptoms, history, and treatments of the patient. The risk of a breach in patient privacy may be estimated heuristically as:

$$R(P) = \left(\frac{a}{t}\right)^{k} \tag{1}$$

where *R(P)* is the risk to privacy, *a* is the total number of terms carried by the agent or other search mechanism, *t* is the total number of terms that describe the current patient, and *k* is the mean commonality of the terms which may be calculated from:

$$k = \frac{m}{p} + \varepsilon\,. \tag{2}$$

The commonality of a medical term is calculated by determining the percentage of total patients, *p*, that are correlated or described by the collection of terms in *a*, *m*, and adding a small error term, ε, to account for the heuristic nature of the equation. Therefore, as *a* increases then so should *m*. Thus, if all terms that describe a set of patients is used, as with current web-based search engines, then the risk of a patient privacy breach is 100 percent. Using only a single term, as is the case with the recommended internet security measures for agent systems, the risk of a security/privacy breach remains small, unless the single medical term is unusual and describes only a very small subset of the physician's existing patients. Future research will examine refinements to the risk equation including changes to the determination of term commonality to better account for very large patient populations, such as those encountered at a state or VA hospital.

The described security protocols and methods have been simulated in the MIRRORS (Medical Information Retrieval Robots and Office Reminder System) agent based medical information retrieval system. Five WWW sites were searched with agents/robots/spiders using medical terms retrieved from a simulated electronic medical record for a patient. The MIRRORS

agent system effectively retrieved 5 documents that were 100 percent relative to treatment of the patient compared to 117 links returned from an existing meta-search engine with an estimated 15 percent relevance, thus information overload is decreased, relevancy is increased, and the most current information is made available.

Assuming that at least one agent was intercepted and assuming that the standard encryption was broken, a medical expert determined that no patient confidentiality would have been compromised. For the case used in the above example, a patient was involved in a motor vehicle accident and suffered probable head and neck injury with the traditional symptoms and signs of head and neck injury. In the best case, the intruder/sniffer will not know the exact office or practice that the acquired information is coming from and in this case the security risk from a single agent's capture is normally less than 2 percent. Risk is defined for the case of multiple, mobile, medical information retrieval (IR) agents as the probability of identifying an individual patient (or very small group of patients) from any information captured from one or more of the mobile IR agents. In the worst case, the intruder will know the exact medical practice that information has been acquired from, then the security risk is estimated at between four and five percent for smaller practices (20-50 patients). Of course, if multiple agents are intercepted or sniffed, then the security risk increases until it is at 100 percent when all of the agents (and correspondingly all terms) have been intercepted.

Conclusion

The abundance and rapid dissemination of information on the internet makes access to this knowledge source a critical business need for medical practitioners (as well as most other businesses in today's e-economy). Agents are becoming ever more present in medical applications and provide a means for assuring accurate and easy to use access to this electronic medium for information. Having access to the most recent medical knowledge will enable physicians to deliver the highest quality of care through better diagnosis, treatment, and expected recovery of their patients. However, accessing medical information (or other (e.g., financial) information) poses security risks, especially in the area of legally mandated patient privacy. Existing methods of encryption employed for internet telecommunications may not be sufficient to guarantee electronic information security. Several techniques for improving on existing security measures have been described and simulated.

Maintaining patient identification information only on a local computer that has adequate physical and access security (e.g., firewalls and password protected entry) is a critical step in maintaining patient privacy security for related medical information that must be transmitted over the internet. Next, non-local or mobile agents that must carry medical information over the internet should be sent to a small number of sources and should only carry a very small portion of the total information related to an individual patient. In this way, even if an agent or subset of agents is intercepted and the encryption is broken, the likelihood of a patient privacy compromise is minimized. A heuristic measure of the security risk from internet eavesdropping has been presented, which accounts for distributing information across multiple internet agents as opposed to a single web search engine. One future research direction is the refinement of this heuristic measure to more accurately predict the information compromise risk associated with mobile agents operating on highly distributed networks. Similar techniques may be developed from the MIRRORS information security methods for other information sensitive domains, such as the legal or accounting domains.

Additional future research needs to investigate the probability that inter-system agents may need to cooperate for effective information retrieval in the future. A protocol for identification of trustworthy agents, perhaps some method similar to digital certificates, must be established and the process of exchanging information should not increase the security risk of a patient privacy compromise.

References

Atkins, D., Buis, P., Hare, C., Kelley, R., Nachenberg, C., Nelson, A. B., Phillips, P., Ritchey, T., Sheldon, T., and Snyder, J. (1997), *Internet Security Professional Reference – 2nd Edition*, New Riders, Indianapolis, IN.

Baujard, O., Baujard, V., Aurel, S., Boyer, C., and Appel, R. D. (1998), Trends in medical information retreival on Internet, *Computers in Biology and Medicine*, Vol 28 No 5, pp. 589-601.

Bazzoli, F. (2000), Gateways to the Internet, *internet health care magazine*, Vol 1 No march/april, 70-77.

Bigus, J. P. and Bigus, J. (1998), *Constructing Intelligent Agents with Java*, Wiley, New York, 1998.

Della Mea, V., Roberto, V., Conti, A., di Gaspero, L., and Beltrami, C. A. (1999), Internet agents for telemedicine services, *Medical Informatics and the Internet in Medicine*, Vol 24 No 3, pp. 181-188.

Detmer, W. M. and Shortliffe, E. H. (1997), Using the Internet to Improve Knowledge Diffusion in Medicine, *Communications of the ACM*, Vol 40 No 8, pp. 101-108.

Finin, T., Labrou, Y., and Mayfield, J. (1997), KQML as an Agent Communication Language, in Bradshaw, J. M. (Ed.), *Software Agents*, MIT Press, Cambridge, MA, pp. 291-316.

Garfinkel, S. and Spafford, G. (1996), *Practical Unix and Internet Security*, O'Reilly & Associates, Bonn, Germany.

Hsiao, D. K., Kerr, D. S., and Madnick, S. E. (1979), *Computer Security*, Academic Press, New York.

Ip, R. W. L., Lau, H. C. W., and Chan, F. T. S. (2000), An intelligent internet information delivery system to evaluate site preferences, *Expert Systems with Applications*, Vol 18 No 1, pp. 33-42.

Maes, P. (1994), Agents that Reduce Work and Information Overload, *Communications of the ACM*, Vol 37 No 7, pp. 31-40,146.

Menczer, F. and Monge, A. E. (1999), Scalable Web Search by Adaptive Online Agents: An InfoSpiders Case Study, in M. Klusch (Ed.), *Intelligent Information Agents*, Springer, Berlin, pp. 323-347.

Rindfleisch, T. C. (1997), Privacy, Information Technology, and Health Care, *Communications of the ACM*, Vol 40 No 8, pp. 93-100.

Silverman, B. G. (1998), The Role of Web Agents in Medical Knowledge Management, *M.D. Computing*, Vol 15 No 4, pp. 221-231.

Silverman, B. G., Andonyadis, C., and Morales, A. (1998), Web-based health care agents; the case of reminders and todos, too (R2Do2), *AI in Medicine*, Vol 14 No 3, pp. 295-316.

Sayeed, L., Hightower, R.T., and Warkentin, M. (1999), An Exploration of a World Wide Web Based Conference System for Supporting Virtual Teams Engaged in Asynchronous Collaborative Tasks, *Proceedings of the 1999 International Conference of the Decision Sciences Institute*, pp. 864-866.

Spath, P. (2000), Case Management Making the Case for Information Systems, *M.D. Computing*, Vol 17 No 3, pp. 40-44.

Tu, H. and Hsiang, J. (2000), An architecture and category knowledge for intelligent information retrieval agents, *Decision Support Systems*, Vol 28 No 3, pp. 255-268.

Walczak, S. (2000), Improving Clinical Care with MIDGI-A, *AMCIS 2000 Proceedings of the Americas Conference on Information Systems*, pp. 314-316.

Zwart, D. and Resnick, H. (2000), The 10 Things Every Training Manager Should Know About TKMä, White paper from Generation21 Learning Systems (available at http://www.gen21.com/nr/Tkm.pdf), Golden, CO.