NETWORK INFORMATION SECURITY IN A PHASE III INTEGRATED ACADEMIC INFORMATION MANAGEMENT SYSTEM (IAIMS)

Steven Shea, M.D.[1,2,5], Soumitra Sengupta, Ph.D.[1,3],
Alan Crosswell[4], and Paul D. Clayton, Ph.D.[1]

From [1]the Center for Medical Informatics, the Departments of [2]Medicine and [3]Computer Science, and [4]Columbia University Academic Information Systems, Columbia University, and [5]The Presbyterian Hospital, New York, NY. Supported in part by a grant from the Digital Equipment Corporation and grant LM04419 from the National Library of Medicine.

ABSTRACT

The developing Integrated Academic Information System (IAIMS) at Columbia-Presbyterian Medical Center provides data sharing links between two separate corporate entities, namely Columbia University Medical School and The Presbyterian Hospital, using a network-based architecture. Multiple database servers with heterogeneous user authentication protocols are linked to this network. "One-stop information shopping" implies one log-on procedure per session, not separate log-on and log-off procedures for each server or application used during a session. These circumstances provide challenges at the policy and technical levels to data security at the network level and insuring smooth information access for end users of these network-based services. Five activities being conducted as part of our security project are described: (1) policy development; (2) an authentication server for the network; (3) Kerberos as a tool for providing mutual authentication, encryption, and time stamping of authentication messages; (4) a prototype interface using Kerberos services to authenticate users accessing a network database server; and (5) a Kerberized electronic signature.

INTRODUCTION

Integrated Academic Information Management System (IAIMS) is a major initiative of the National Library of Medicine whose goal is to sponsor research and demonstration projects at the institutional level that provide integration of computing systems in academic medical centers.[1] In the context of IAIMS, integration involves cooperation and connectivity of computing infrastructure over several domains, including hospital information systems, library resources, administrative systems in medical schools and their affiliated hospitals, and databases assembled for research, quality assurance, and other non-routine purposes. Integrative activities occur under IAIMS at several levels, including the planning and design of the institutional architecture for hardware, software, communications, and interfaces, as well as the data and knowledge representations in these computing environments. IAIMS is an effort to apply the concepts and achieve the benefits of enterprise computing in the setting of academic medical centers.

In 1988 Columbia-Presbyterian Medical Center received funding for Phase III IAIMS.[2] A great deal of progress has been made over the last four years in implementing a full scale IAIMS. Some of the landmarks in this effort include installation of a campus-wide network that reaches 13 buildings, including all major buildings on the 168th Street campus and the Allen Pavilion, a 400 bed hospital located three miles to the north. This network is heterogenous and comprises major segments of token ring linked via gateway to additional major segments of ethernet. Approximately 500 personal computers connected to this network have been installed under IAIMS, and there are approximately 200 additional personal computers that users own and that have been connected. In addition, there are many larger computers that offer applications and databases to network users, including two IBM 3090 mainframe computers and more than 30 DEC VAXes, three HP 3000s, one Data General 15000, one IBM AS/400, and multiple Unix-based minicomputers such as DEC 5000, DEC 3100, IBM RS/6000, and AT&T 3B2. Major databases and applications that have been made available on this network include the Clinical Information System (laboratory results, radiology reports, discharge summaries, pathology reports, demographic profiles of patients, clinical profiles of ambulatory patients, and other clinical data), the Library Information System (Medline, CLIO - an on-line catalogue of library holdings - and several locally mounted documents such as the syllabus of the medical school's Abnormal Human Biology course), and electronic mail. A workstation that will manage concurrent access to multiple applications and databases resident on different machines and that will provide capabilities to move data across these applications has been designed and prototyped.

A major challenge confronting all information system developers and managers in health care is control of access and protection of data integrity and security. The need to address this issue in the context of IAIMS was identified by members of the Columbia-Presbyterian IAIMS External Advisory Committee, who suggested that external consultants be retained. Traditional approaches to data security have involved access control at the level of the user, machine, ap-

0195-4210/92/$5.00

plication, which operations are permitted, and which parts of data sets can be accessed.[3] This model fails in several critical ways in the IAIMS environment developed at our institution. It is not practical to prohibit access to most medical data for most patients, since the most sensitive data is often the most important. It is critical that nurses and houseofficers know the medical history as well as (for example) the blood type of all patients who may come to the emergency room or have a major complication in the hospital. Thus a softer edged philosophy has been accepted in many medical centers with well developed data systems, and the audit trail has become a critical tool for managing data security issues. A second problem is that the databases comprised by IAIMS are distributed among many applications and multiple heterogeneous machines. Users cannot easily log on and off each server when switching applications during a session. Third, machine- or application-based security does not address security at the network level. Thus, management of access, authentication, and data security has become a critical integrative function in IAIMS and a network-wide task that bridges many machines from multiple vendors, each with its own security protocols.

SPECIFIC PROJECTS

1. Policy Development

During the spring and summer of 1990, the Medical Center together with Columbia University retained Kroll Associates as consultants and initiated a review of data security needs. Approximately 80 people participated in twelve committees. Committee members included senior officials and legal staff of the Hospital and University, physicians, faculty, library personnel, systems managers, and technical staff. Agendas were developed by the consultants, the chair, and the recorder prior to the meetings.

Each commitee submitted a report with recommendations. At a summary meeting of committee chairs and recorders and the consultants, these reports were presented. The major conclusions were that:

* In general, users were unaware of the need for security and valued ease of access over security.

* There had, in general, been a lack of dialogue about computer security both within the Hospital and the University and between counterparts in these two institutions, that the dialogue created by the security review process had uncovered many common areas of concern, and that the basis in shared interests was present for working together to address security needs.

* System managers felt a lack of strong commitment from senior officials for strong security, particularly in the University environment where freedom of information was highly valued.

* There were important gaps in the policies and procedures in both the Hospital and the University with regard to computer and data security.

The recommendation given the highest priority was to address the need for additional data security policies and procedures. These policies and procedures, once articulated, would provide a framework within which risk analyses could be conducted and technical solutions could be implemented.

As a direct result, a Joint Columbia University-Presbyterian Hospital Data Security Committee was formed and has been meeting for the last year. This committee has high level representation from both the Hospital and University and includes people in administrative, academic, and clinical areas. Policies are being drafted to address such areas as maintenance of lists of authorized network users, requirements for transfer of data files among servers ("downloading"), use of unlicensed software, tampering with system hardware or software, and snooping.

2. Authentication Server

In order to provide network-wide security, the Medical Center needs centrally managed authentication and authorization services. We have installed a DECStation 5000/200 running UNIX as our prototype network authentication server. It uses MIT Kerberos protocols for authentication (see below) and maintains the tables of data for users ("handles" [primary key], name, status, passwords, privileges, etc.) and services for authorization information. These data are maintained using a relational data model. A key concept is the **handle**, which has the following characteristics:

Handles are unique.

Handles cannot be assigned to more than one user or reused. Once a handle has been issued, it can never be re-issued even if the original user has left the University and/or Hospital.

Handles are stable. Handles therefore cannot contain information identifying a class of user (e.g., nurse, student, physician), since users may change their class characteristic. Users may also change their passwords; indeed, some systems require periodic changes in passwords. The handle therefore cannot be the password, since the handle must remain stable.

Each user can have only one handle in the authentication server, mapped to multiple passwords (one for each client), if necessary,

The X.500 equivalent to the handle is the "distinguished

name", except that X.500 does not include the contraint that the "distinguished name" cannot change when the user class characteristic changes.

A working prototype database has been created and loaded with approximately 15K records for University users. The next step will be to load data for Hospital users, thereby demonstrating ability to reduce disparate user formats to a common data structure in the authentication server database. A task to be completed during the next year is to create working interfaces to University and Hospital systems such as Personnel and Payroll to provide immediate updates and purges from the authentication server database. Our direction is for this database application to evolve into an ISO X.500-based services directory.[4]

3. Kerberos

We are using Kerberos [5] as a tool to address several key requirements for network security:

Encryption of authentication messages (passwords). Kerberos encryption conforms to the Data Encryption Standard (DES).

Time stamping of authentication messages. Time stamping increases the security of authentication transactions by invalidating purloined passwords, encrypted or not, after a short period of time. Kerberos provides time stamped tickets for access to network-based servers.

Mutual authentication. A server receiving an access request from a user should be able to authenticate the user. In a distributed environment, it is also important for the user to know that he/she is communicating with the server, and not an imposter. Thus mutual authentication is required. Kerberos supports mutual authentication. The secret (password) is shared by both the user and the server with a trusted third party (the authentication server/Kerberos key distribution facility).

It follows that all servers on the network must support Kerberos or a Kerberos-like facility. In principle Kerberos is part of the OSF Distributed Computing Environment (DCE), but working Kerberos libraries are not in fact available for some hardware/software platforms in our environment. The necessary interfaces will be built during the next two years.

4. Interface Prototype

A "session manager" program using Kerberos services on the authentication server is currently in production use, supporting user access to multiple computing services across Internet and other local networks in a transparent way. The design of the "session manager" is based on an "application gateway" concept, in which a user manipulates his/her sessions with many applications through a single application gateway computer. This approach permits the network to be conceptually partitioned into "secured" (gateway to server) and "non-secured" (client to gateway) categories. Kerberos provides authentication services for the "non-secured" end. In addition to authenticating a client user, the system permits only the specific kinds of access that a user is allowed for each application (as specified in the database). The first server to follow the Kerberos protocols uses the DECStation 5000 as the application gateway to MEDLINE running on our IBM mainframe in the Hospital's data center 20 miles from the Morningside campus, where the users are located. In this example, users access the application through a local area network which is relatively vulnerable to illegitimate access, while communication between the application and the mainframe is through an SNA network, which is more difficult to break into.

5. Electronic Signature

Our Clinical Information System maintains and displays to users on the network a large number of dictated/transcribed free-text reports that require an authenticated signature on their final forms. These documents include discharge summaries, radiology reports, operative reports, pathology reports, and attestations to discharge diagnoses. Presently these documents are signed on paper in the Medical Records office. The electronic signature capability will make it posssible for physicians to sign these documents from any networked workstation. As the database for these free-text documents we are migrating from an IBM AS400-based system to a RISC/UNIX base open platform with a free-text editor and a document manager that supports the Revision Control System (RCS) for version management. A working prototype of the electronic signature has been been built. Features include use of Kerberos services on the authentication server for encryption, time stamping, and mutual authentication; interface to our electronic mail system so that users signing documents can e-mail copies to colleagues, such as referring physicians; and listing of unsigned documents, also interfaced to the e-mail system so that e-mail users can be notified of documents needing their signature. Electronic signatures will also be used for administrative purposes, such as purchasing and personnel action, once these forms are available electronically.

ADDITIONAL TASKS

The security project has defined several additional objectives that will be addressed over the next two years. First, it is necessary to achieve control of the flow of newly authenticated users into the system and removal from the system of users who are no longer current. This will require close coordination with personnel and student database managers in the Hospital and University and the creation of interfaces to these databases. Procedures will be established to

automate incremental additions and deletions to resource definitions in the authentication server. Existing local security systems will, in general, require modifications to accept authentication messages from the authentication server. It will be necessary to develop tools to integrate the authentication server databases with existing security systems. It will also be necessary to develop standards for new systems on the network so that connectivity with the authentication server can be considered in the design of new systems. The institutional security policy process described earlier will be critically important in creating willingness among local resource administrators to accept these modifications.

Second, we will build tools to monitor and analyze access using audit trails. Two methods for analyzing audit trails have been proposed, namely use of expert systems and use of statistical methods.[3] Problems we face include identification or development of tools for an expert system, such as an editor whereby rules can be written, a database in which to store the rules, an inference engine, a report generator, and a monitor to track triggering events. It will be necessary to store the audit trail in a form such that the inference engine can read the contents. For example, in dealing with clinical data, the data required by the inference engine will include characteristics of the patient (patient is male or patient is 50 years old) and characteristics of the system user (the user is an obstetrician or a pediatrician), such that a rule (e.g., in general, obstetricians should not be looking at clinical data of male patients, or pediatricians should not be looking at clinical data of adult patients) can be triggered. Statistical approaches may also offer solutions to the problem of processing audit information. Here the main challenge is to derive probabilistic rules from data access patterns of user classes. In areas of the IAIMS environment where classes of users are prevented from accessing certain classes of data, e.g., secretaries accessing faculty salary databases, audits should be maintained of system rejections of such read attempts. It will also be necessary to develop sampling strategies in order to protect system performance. These strategies will require that the database include characteristics of the data classes. For example, the sampling fraction may be higher for HIV tests or salary data than for routine urinalyses or purchasing data. The data collected through audit trails will also need to be compiled into summary utilization reports and summary reports of audit and rule triggering events. Local database administrators will need to be consulted regarding levels of sensitivity of data within their own domains.

Third, we will develop tools for continuous performance monitoring and analysis of the authentication server.

Finally, we need to create a production environment for the security server. Because of its functions, this server represents a single point of failure for all network services requiring user authentication for access. In addition to being highly secure, it must have robust recovery facilities and redundant network access paths, storage subsystems, and server processor availability in the event of hardware failure.

CONCLUSION

Security is a critical aspect of the management of distributed computing systems. Network-based security is a new concept, with few models on which to draw in the medical field. A major need identified through an intensive institutional review was for comprehensive data security policies that span all users, machines, applications, and data that are interconnected by IAIMS or other enterprise-wide connectivity initiatives. Kerberos promises a platform on which useful implementation experiments can be conducted. IAIMS provides a context in which our institution has made commitments to address the security requirements of such a distributed system.

REFERENCES

1. Improving health professionals' access to information: Challenges and opportunities for the National Library of Medicine. Report of the Outreach Planning Panel to the Board of Regents of the National Library of Medicine. 1989.

2. Anderson RK, Clayton PD, et al. Integrated academic information system (IAIMS) implementation at Columbia-Presbyterian Medical Center. New York: Columbia-Presbyterian Medical Center, 1988.

3. National Research Council. Computers at risk. Safe computing in the information age. Washington, D.C.: National Academy Press, 1990.

4. International Organization for Standardization. Information processing systems - open systems interconnection - the directory. DIS9594. (CCITT X.500 Recommendations). April, 1988.

5. Miller SP, Neuman BC, Schiller JI, Saltzer JH. Kerberos authentication and authorization system. Section E.2.1. Project Athena Technical Plan. MIT, 1988.