MEDINFO 95 Proceedings. R.A. Greenes et al. (editors)

Security Management in the Healthcare Environment

Furnell S. M.[a] and Sanders P. W.[a]

[a]*Network Research Group, Faculty of Technology, University of Plymouth, United Kingdom*

Modern healthcare establishments increasingly rely on information systems in all aspects of work; any compromise of their security may represent a significant threat to both the organization and the patient. This paper discusses the increasing need for standardized levels of protection in healthcare computing systems and networks, outlining steps that have been taken to achieve this within European establishments. The paper then considers specific technical concepts that may be applied to improve security in healthcare at both local and international levels.

1. Introduction

As with many other areas of society, the healthcare field has been significantly affected by the adoption of information technology. Modern establishments now utilize a wide variety of equipment, ranging from stand-alone PCs to minicomputer or mainframe systems; they represent significant assets of the business. In addition, many organizations now incorporate links to remote sites via Wide Area Network (WAN) arrangements, with increasing volumes of data transmitted between different establishments. This is likely to increase still further with the proposed standardization of computerized health records using a common data structure [1].

The adoption of information technology has been accompanied by an increase in the number and variety of medical applications, which now affect most areas of operation (including patient care, finance, staffing, administration, etc.). As a result, healthcare professionals have become increasingly dependent upon the availability of computer systems and reliant upon the correctness of the data that they hold.

The above trends highlight an increasing need for security in healthcare systems. Information systems may be compromised by a variety of accidental acts or by deliberate, malicious activity (e.g., hacking, fraud, virus infection, etc.). As such, it is now recognized that security issues must be considered during the design and development of new health information systems. In addition, security must also be added or enhanced in many existing systems, which were originally implemented without such considerations in mind and, consequently, have no standard arrangements.

2. Security Requirements in Healthcare

As with many other application areas, security requirements in healthcare are centered around the issues of confidentiality, integrity, and availability [2]. These may be achieved by incorporating security services for authentication, confidentiality, integrity, and non-repudiation as defined by ISO [3].

The nature of the healthcare environment tends to impose constraints on the types of protection that will be considered acceptable. For example, measures that greatly interfere with users' abilities to perform their primary duties (e.g., care delivery) will not be tolerated. This points to a requirement for measures that are as simple and transparent as possible. In addition, financial cost is an important consideration as investment in security is often hard to justify against expenditure that would improve patient care. As a result, the use of software-based technologies may be a more realistic approach for widespread adoption than expensive hardware-oriented methods.

Despite these constraints, the increased interconnection and the sharing of data between different establishments heightens the need for uniform levels of protection throughout the healthcare community.

3. Baseline Security for Healthcare Systems

The need for improved security is already recognized within Europe and has been addressed by the CEC SEISMED (Secure Environment for Information Systems in MEDicine) project, with which our group has been involved [4]. The objective of SEISMED is to provide practical security advice and guidance to all members of the healthcare commu-

MEDINFO 95 Proceedings. R.A. Greenes et al. (editors)

nity who are involved with the development, operation, and management of information systems.

Part of the project has been dedicated to the development of baseline security standards for existing systems and networks, describing the levels of protection that are appropriate for the healthcare environment. It is envisaged that these will eventually help to form a common reference for the security of healthcare systems within Europe.

The guidelines for existing systems highlight ten key principles of security that must be considered: 1) security policy and administration; 2) physical security; 3) disaster planning and recovery; 4) personnel security; 5) information technology facilities management; 6) user identification and authentication; 7) system access control; 8) database security; 9) system maintenance; and 10) legislation compliance. These principles encompass a very wide range of considerations, with coverage ranging from general security concepts to more specific technical measures. In addition, the networking of medical systems has been recognized as an important issue in its own right. While networks offer significant opportunities for improving healthcare services (thanks to the increased availability and sharing of information), there are also inherent security considerations. Examples of network threats include wiretapping, message replay, message repudiation, and user impersonation. The SEISMED guidelines for networks present a further set of baseline standards to counter these and other threats and are primarily based upon encryption.

The definition of a healthcare baseline represents a significant step in achieving the desired standardization of protection in the field. However, while the baseline standards provide comprehensive guidelines on "what" aspects of security should be considered, they do not attempt to describe in any great detail "how" technical measures may be best implemented. A comprehensive and flexible security system is needed that can be integrated into applications as required. The remaining sections outline how such a system may be realized in the healthcare environment.

4. Use of Trusted Third Party Techniques

To meet the specific network security requirements for both local and wide area systems, a unique and unforgable identification of all potential users (perhaps on a global scale) is necessary. These identities must be authenticated and "binded" to the activity or data used in that session. It may be appropriate to use a naming and registration policy and infrastructure based on the international standards and technical framework of X509/ISO 9545-8 [5]. Non-repudiation of the activities is required, together with confidentiality and data integrity during communication. Most methods to achieve these services are based on secret key cryptography and involve digital signatures, the encryption of data, and the support of Trusted Third Party (TTP) infrastructures for wide scale use.

The implementation of such an arrangement involves public key systems, such as the RSA algorithm, with smart card technology for transparency and ease-of-use by the healthcare staff. The cards perform various cryptographic functions (the creation and verification of signatures, encryption/decryption of data, and storage of secret keys and other sensitive data) and perform other special functions particular to the application. The TTPs act mainly as Certification Authorities for the digital signatures they provide and, while they give a value-added service, must be trustworthy beyond the level of normal computer systems.

In order to provide all the necessary functions on an international scale, a network of TTPs is required, as shown in Figure 1. At this level the infrastructure will be generic for all applications, but at the local domain and sub-domain levels (as shown in Figures 2 and 3), specific operations can be incorporated to satisfy HCE security policies, with the TTP being extended to a more comprehensive Security Management Centre (SMC) set of functionalities [6].

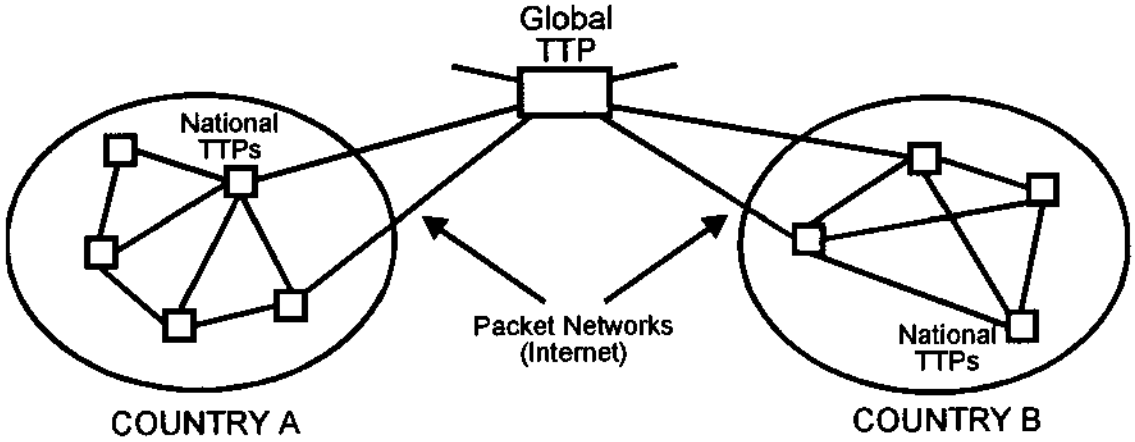

Figure 1. TTP infrastructure at an international level

MEDINFO 95 Proceedings. R.A. Greenes et al. (editors)

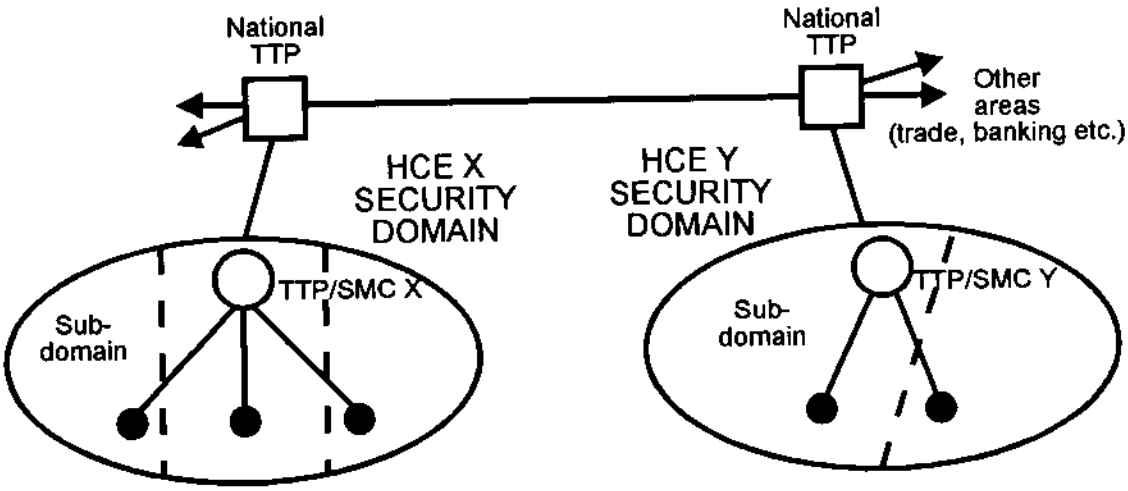

Figure 2. TTP infrastructure at a national level

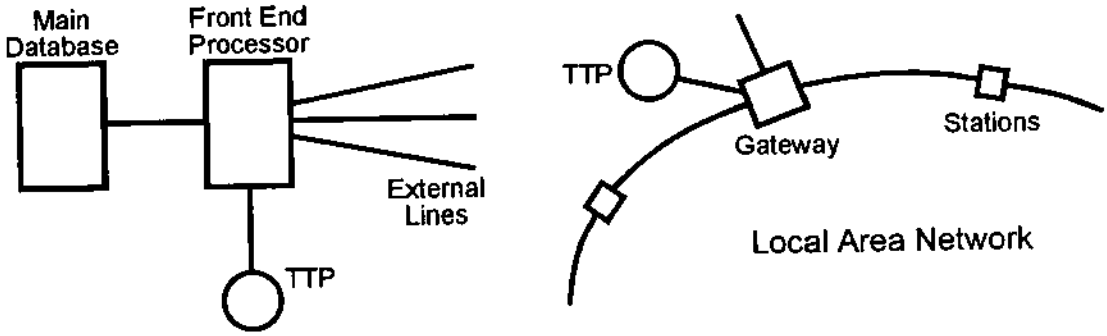

Figure 3. TTP infrastructure at a local level

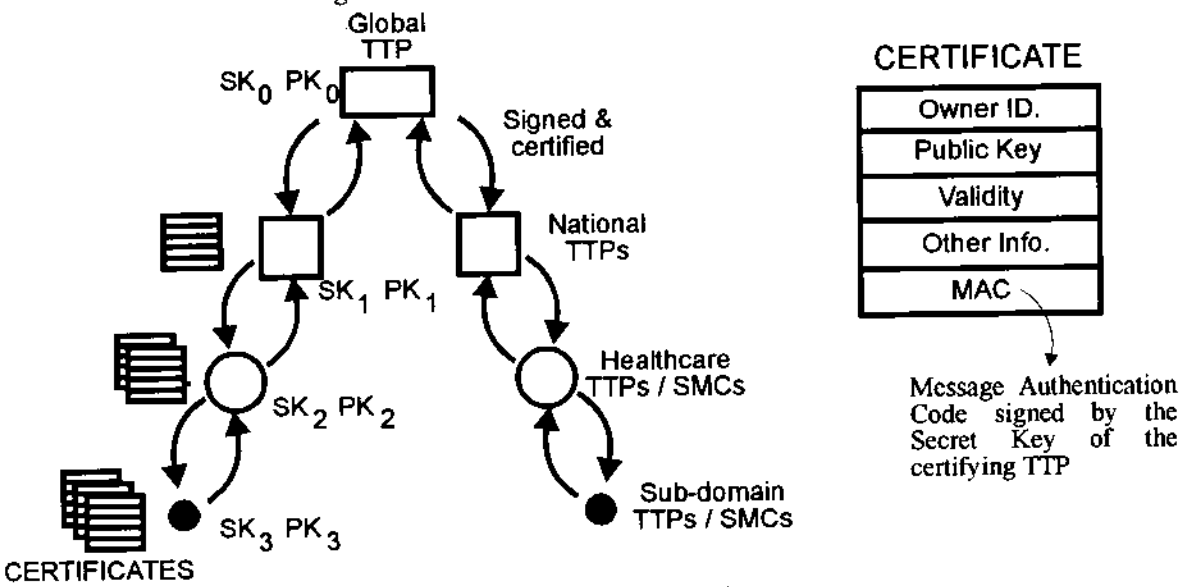

Figure 4. Logical TTP hierarchy

In order to guarantee the authenticity of certificates, a hierarchical certification structure is used. This is shown in Figure 4, along with the format of the certificates, illustrating how additional certification occurs at each TTP level (with SK and PK representing the secret and public keys of the TTPs at each stage).

MEDINFO 95 Proceedings. R.A. Greenes et al. (editors)

Each TTP in the hierarchy is certified by the TTP in the next layer up, which not only provides the credibility of the complete system by defining the individual certification path within a certificate, but allows for the loss of a hierarchical level under fault conditions (with the next higher order certificate being used). The arrangement is common to the X509 Directory services architecture.

The TTP network can provide or verify signatures via the certificates, facilitating authentication and non-repudiation services. In addition, secret keys can be passed between users in a hybrid system where a symmetrical algorithm is used to provide confidentiality. Finally, the integrity of data can be confirmed by the signing of a Message Authentication Code that is a hash function of the message.

As previously mentioned, additional security services can be incorporated into the TTP overlay in the local HCE security domains. This is discussed in the next section.

5. Real-Time Supervision

While the TTP will ensure the integrity and confidentiality of operations, an additional mechanism may be required within the local HCE domains to ensure that users are continually authenticated during their session and that they do not act outside their permitted bounds. A solution is to incorporate a real-time supervision system to detect the unauthorized activity and strengthen standard authentication and access controls.

The supervisor would use expert system techniques to compare user and process activities within the domain against models of normal and suspicious behavior, thus revealing any potential security problems (i.e., if an activity is incompatible with normal behavior or is compatible with suspicious behavior then it may be an intrusion). These models may be represented by maintaining behavioral profiles (for normal activity) and using pre-determined intrusion indicators (for suspicious activity).

It is considered that behavior profiling may operate at two levels. At a high level it is possible to classify users according to their role within the HCE, developing general rules for acceptable activities within each class. In addition, lower level profiles can be developed for individual users by analyzing their use of the system. Measurable characteristics may include application and file usage, typical access times and locations, individual keystroke/typing patterns, and instances of login failures or access violations. Validation of activity against the high level profile should ensure that users are operating within their legitimate bounds, while the lower level also allows authentication of the subject according to the behavioral characteristics. The user-specific profiles would need to be refined over time to account for legitimate changes in behavior.

In addition to using profiles, the supervisor would monitor the system at a more general level to identify suspicious activities that may form part of a compromise attempt. Examples of such indicators may include access of infrequently used files, consecutive access violations, and extensive/frequent use of "help" systems. While none of these events alone would be conclusive of an intrusion, they could be used as a trigger for a more detailed monitoring or investigation. The disadvantage of this approach is that it will only cope with known intrusion scenarios.

Supervision could operate continuously throughout a session or at random periods, depending upon factors such as system load and application sensitivity. In either case, it would operate transparently, unless an intrusion was suspected (in which case the system manager would be alerted and/or other appropriate safeguards would be taken).

The implementation of supervision in this manner is compatible with the desire for a software oriented approach to security as described in section 2.

6. Conclusion

A European-wide network is already operating on a prototype scale [7], with extensions to the HCE being designed at present. It is expected that this approach will provide a relatively cheap and easy-to-use service, facilitating effective security for healthcare establishments.

7. References

[1] European Commission. *Good European Health Record (GEHR)*, EC AIM A2014, EC. Brussels (1993).

[2] Commission of the European Communities. *Information Technology Security Evaluation Criteria (ITSEC).* Provisional Harmonized Criteria, CEC, Brussels (1991).

[3] International Standards Organization (ISO). *Information Processing Systems* - OSI RM. Part 2: Security Architecture, ISO / TC 97 7498-2 (1988).

[4] European Commission. *A Secure Environment for Information Systems in Medicine (SEISMED).* EC AIM A2033, EC, Brussels (1991).

[5] International Standards Organization (ISO). *Information Technology - Open Systems Interconnection - The Directory.* Part 8 : Authentication Framework, ISO / IEC 9595-8 (1990).

[6] Muftic S., Patel A., Sanders P., Colon R., Heijnsdijk J., & Pulkkinen U. *Security Architecture for Open Distributed Systems.* Wiley Professional Computing (1993).

[7] Computer Security Technologies (COST). *Smart Card System Specification.* Sweden: Hasselby (1994).