Security Issues on the Internet*

Josef Pieprzyk †and Andrzej Goscinski ‡

Abstract

The paper presents an overview of security issues in the Internet. First security threats are discussed. Next some general issues of access control are considered with firewalls as a special case of the access control which allow to separate parts of the internet using packets screening techniques. Two solutions of firewals are shown, one based on so-called cryptographic tickets and the other based on a network security probe. Later a discussion on intrusion detection systems follows. The last part of the paper shows how to distribute cryptographic keys via insecure channels (Internet).

1 Introduction

A rapid advancement in computer, communication and information technologies has brought us new exciting information services. Most of these services are remotely accessible via efficient, reliable and cheap communication networks. Being in Australia, persons can be in contact with their friends in the North America, Europe or Asia in almost the same way as they communicate within their own organization. The first time in the human history, interest groups and their activities are not limited by physical distances - the only real boundary is the human imagination.

As the range of computer and communication services is growing, the interactions among people who have access to international computer network become richer. Unfortunately, there is also darker side of the proliferation of information technology and communication services - an unauthorized or illegal access to information also called a computer crime. At the beginning of the "information explosion" age, the legal system was not prepared to deal with computer crimes. Now all industrially advanced countries have the law to combat the abuse and misuse of information services. But there is a lot to be done - there is no international convention which would deal with computer crimes committed in one country while the culprit sits against their terminal is the other country.

Apart from legal issues, there are many technical problems related to the secure computations and communication. In general, the designers of secure services (systems) would like to allow the access to all authorized users who act ethically (according to their access rights). At the same time the service (the system) should become unworkable if either users are illegal or they exceed their access privileges.

This work addresses the problem of secure information flow through the computer networks.

*presented at IT Security Confrence, IIR, Canberra, April 10 & 11, 1995

†Centre for Computer Security Research, Department of Computer Science, University of Wollongong, Wollongong, NSW 2500, AUSTRALIA, e-mail: josef@cs.uow.edu.au

‡School of Computing and Mathematics, Deakin University, Geelong, Victoria 3217, AUSTRALIA, e-mail: ang@deakin.edu.au

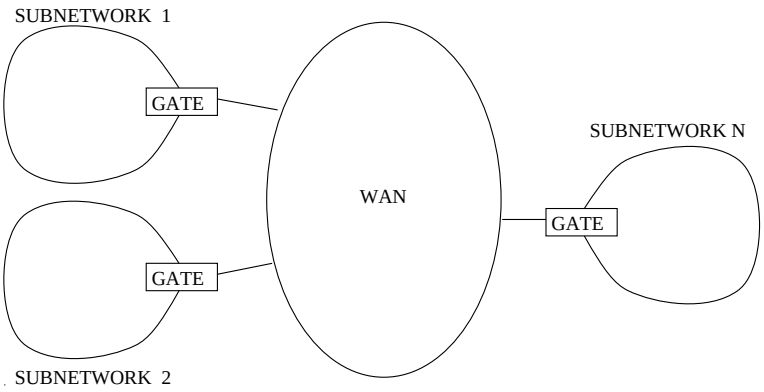

Figure 1: General Internet structure

2 Security Threats in the Internet

The general structure of the Internet is shown in Figure 1. Subnetworks can be identified with local area networks (LAN) owned by institutions or organizations. The wide area network (WAN) connects all subnetworks together so we can think about the whole system as a single network often called the Internet. LANs are viewed as communication networks that connect data communication devices within small areas. Sometimes it is convenient to identify a LAN with a communication network which connects a collection of workstations and servers which are managed by a single (distributed) operating system (see [5]). On the other hand, WANs usually cover large geographic area (city, state, country, region, world) and may consists of many subnetworks - this causes that WANs may have a complex hierarchical structure with complex management system. The data flow from a local subnetwork (LAN) goes to the WAN via a gate. Gates provide interfaces so the local computers can talk to the other computers in the Internet. Also gates can be used to force some specific security policies related to the incoming data.

The world wide network (Internet) is a complex systems which covers most of countries. Because of its size and structure, the Internet is vulnerable to many computer security threats. By their nature, it is impossible to physically protect communication channels. For example radio channels are wide open for potential eavesdroppers. Also it is not difficult to imagine that an enemy controls a part of Internet so they cannot only read the transmitted information but also inject fraudulent messages (packets) into the network.

In general threats to the security of the Internet can be classified into the following classes ([10]):

Confidentiality threats: they are related to disclosure of information to unauthorized parties (for example a release of the confidential medical record of a patient).

Integrity threats: they are centered around modification, replay or injection of information by illegal parties (for example modification of the amount of money in the ATM transaction).

Availability threats: they are related to denial of service by part of the Internet (for example damage of the gate causes that all users of the local LAN cannot use the Internet and also the Internet parties cannot communicate with them).

Accountability threats: they concentrate around threats to the correctness of authorization, audit, and monitoring systems.

	O_1	O_2	O_3
S_1	-	execute	read, write
S_2	wait	signal	read

Table 1: Example of access matrix

To protect the Internet (and its components) against security threats, it is necessary to address the four issues that correspond to the threats classes mentioned above.

Confidentiality is the prevention of unauthorized disclosure of information. In any information system confidentiality is forced by an appropriate security mechanism (which follows the confidentiality policy). As the confidentiality has to be imposed in different parts of the Internet, the mechanisms depend upon the place. In parts where the physical security can be forced (by a reliable identification), an access control mechanism is used. In some other parts such as communication channels where the physical security cannot be imposed, cryptographic techniques are used to make the information flow unintelligible. If an enemy is expected to access communication channels only, encryption is applied between the Internet nodes (hosts) - this is called the link or host-to-host encryption. If however the enemy can penetrate a part of the Internet, encryption is used between two communicating parties - this is the end-to-end encryption (the sender encrypts and the receiver decrypts - the whole information is transmitted via Internet as a ciphertext stream).

Integrity is the protection against unauthorized modification. Depending on the place of the Internet, it is possible to talk about the packet integrity (inside the WAN), the message integrity (inside the local host), or the message stream integrity (inside the user workstation/LAN). The integrity protection can be implemented differently depending on the security policy. These variations can be: sender authentication, sender non repudiation or receiver non repudiation.

Availability is the protection against accidental or malicious activity. Wherever it is possible the most sensitive parts of the Internet should be physically protected. There should be a recovery strategy in place for each organization (LAN). The communication networks should be designed taking into account robustness and fault tolerance of the Internet. The traffic in the Internet should be constantly monitored so if someone wants to block the communication (by saturation of a part of the communication network), the part of the Internet with illegal activity can be quickly identified and contained.

Accountability is the proper identification of responsibility for an action to the correct party in the Internet. Accountability includes all methods and techniques to ensure that the security mechanisms are working correctly and are unpenetrable for illegal parties.

3 Access Control

Access control is widely used to ensure that the active entities (such as processes, users) access computer resources according to the accepted protection policy. The active entities are called *subjects* and the passive entities (resources) - *objects*. As any subject has a specified set of access rights to any object (note that if the set is empty, the subject cannot access the object), it is natural to illustrate the current state of access privileges by so called *access matrix*. The rows of an access matrix are assigned to the subjects and columns - to the objects (see Table 1). The entry (S_i, O_j) specifies how the subject S_i can access the object O_j by enumerating the suitable access privileges. The collection of possible privileges depends on the type of object. For files, a

typical set of access privileges is: read, write, create, delete. For processes, execute, wait, signal, send, receive are possible access rights.

Mandatory access control. Entries of an access matrix are determined exclusively by the system (users cannot interfere with the matrix). Access rights are or denied to a user depending on the classification of object and the clearance of the subject (user).

Discretionary access control. Entries of the access matrix are determined by users who "own" objects. It is left to the owners' discretion to grant, modify or deny access for other users. A typical set of modification commands includes: transfer of an access right to another subject, grant an access right, delete an access right, create an object, delete an object, create a subject, delete a subject.

For the mandatory access control, the access matrix approach is not adequate because it does not take into consideration the semantics of the information. The lattice security model is an extension to the access matrix model to include classifications, clearance, and rules concerning the classifications (for more details see [5]).

For the discretionary access control, the access matrix becomes dynamic. Changes are being introduced by the owners using modification commands. There is, however, a general problem of security verification of the access control assuming dynamic access matrix with modification commands. It was proved ([4]) that the problem of security verification of such a system is undecidable.

Access control provides a balance between sharing (computer /network resources) and protection. Access control granulation depends on the level of sharing:

- no sharing (complete isolation),
- sharing objects,
- sharing untrusted subjects.

No sharing usually provides perfect security but is extremely inflexible and not practical. Better granulation, flexibility and usage of computer or network resources is provided by access control with shared objects (most of the existing operating systems). Sharing untrusted subjects provides very flexible access control with sometimes poor security (as there are practical problems with the security verification).

The Internet with its resources scattered around the world generates a new challenge for the access control. Most of already existing access control implementations, i.e., methods, techniques and technologies, have been developed for operating systems located in a single computer. However, there are two methods, used in experimental distributed operating systems, which could be employed to protect resources of the Internet: access control lists and capabilities. Recently the concept of isolation of the Internet parts have been introduced. This isolation is commonly referred as *a firewall*. The firewall concept includes methods which can be applied to control flow of packets through internetwork gates so "illegal" packets are identified and the further action is initiated.

In the next section we are going to describe these three access control implementations, and in particular two access control mechanisms to illustrate the concept of firewalls.

4 Implementation of the Access Control

4.1 Access Control Lists

The access control list is a scheme which is a result of decomposing the matrix by columns. The resulting list for each object consists of ordered pairs

$$< subject, rights >$$

An access list-based system can use either process names or unique identifiers, as subject entities in the pairs.

According to this scheme each object is associated with an object monitor, which is responsible for the control of direct access to this object. Thus, the access monitor maintains and uses the access control list which enumerates the subjects (users) that have access rights to this objects and defines these rights.

The advantages of the access control list scheme is that for a given object it is possible to determine easily and efficiently which subjects have access rights, and to revoke rights of a specified subject. The disadvantage of this scheme is that it is difficult to determine rights of a given subject as they are scattered around all objects.

4.2 Capabilities

The capability is a scheme which is a result of decomposing the matrix by rows. The resulting list for each subject consists of the objects to which the subject has access rights and determines these rights.

$$< object, rights >$$

Thus, a capability is a unforgeable ticket which specifies one or more permission modes and allows its holder (subject) to access a named object.

The advantage of the capability-based scheme is that for a given subject it is easy to determine all objects that can be accessed and the access rights. However, this scheme has some disadvantages. Firstly, for a given object it is difficult to determine which subjects can access it and what access rights they have. Secondly, revocation of privileges is a complicated task, and this operation can leave a number of copies of obsolete capabilities in the system.

A comparison of capability based access control scheme to access lists shows that capabilities support closer adherence to the principle of least privilege. Furthermore, they allow the passing of access rights between subjects without mediation by an object monitor.

4.3 Firewalls

Firewalls provide access control via screening data packets which are transmitted between two parts of the Internet. The access control and the corresponding access matrix deals with subjects and objects which are defined on higher layers of the network (usually on the application layer - for a quick reference to the network architecture, see ISO OSI reference model [5] or [15]). The screening is being done on the network layer. Therefore the main problem in implementation of access control based on firewalls is how to create a secure link between subject (or object) and the packets which correspond to it. This link can later be used to distinguish "friend" packets from "foe" ones.

Network tickets. The solution discussed in [9] describes an access control mechanism which applies cryptographic "tickets" to decide whether the packet is genuine (there is the corresponding subject) or forged (there is no corresponding subject in the protected domain).

For a packet n which is associated with a client (subject) C, the ticket is calculated as

$$ticket_C(n) = DES_{K_C}\{S_C \oplus n\}$$

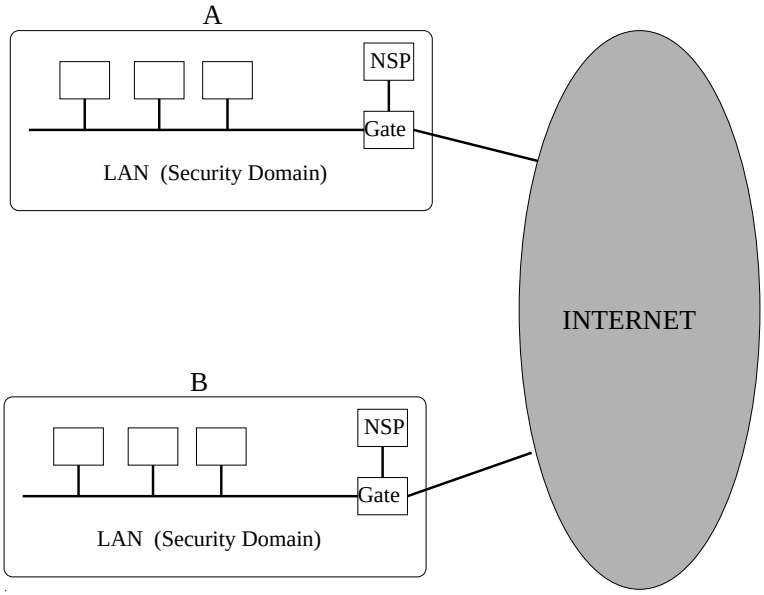

Figure 2: Application of NSP in the Internet

where K_C is a secret cryptographic key which is shared between the client host and the server host. S_C is a unique client identifier/name and n is a sequence number of the packet associated with the client C (or the corresponding session).

The server host computes tickets and attaches them to corresponding packets. The client host takes the information from the packets headers (address of the client which identifies S_C and the sequence number) and recalculates the ticket. If both tickets (attached and recalculated) are not equal, the client host kills the packet.

Note that the ticket depends on the client (subject) identity and the sequence number n - there is no relation between the ticket and the packet data. An illegal party can modify the packet data without any chances that the modification will be detected by the client host. The problem can be solved if the ticket is calculated as

$$ticket_C(n) = DES_{K_C}\{S_C \oplus n \oplus MD5(Data)\}$$

where MD5(Data) is the digest of packet data produced by MD5 hashing algorithm. This time the ticket is also used for the integrity checking.

In general, calculation of tickets by the server and client hosts, introduces a delay which may not be acceptable where the speed is crucial.

Network security probe (NSP). In the previous solution, client hosts treat every access as a dangerous one - this results in heavy overhead and an additional communication delay. It is possible to apply a softer approach to the network access control - all access attempts are legal until the client host gathers enough evidences to conclude that the communication (of packets) is illegal. The solution discussed here was presented by Rolin, Toutain and Gomabault in [11]. NSP can be used to monitor a single security domain (single LAN) or a data flow between two security domains (two LANs A and B - see Figure 2).

NSP observes the communication and after detecting an illegal activity, it can decide to act. The action can be twofold: closing the illegal connection and/or notification of the security officer for further action. The architecture of the NSP is shown in Figure 3. Monitoring is being done by the listener module. The analysis is done by the security analyzer which uses the security policy specification stored in the data base. The analyzer also interacts with specialized

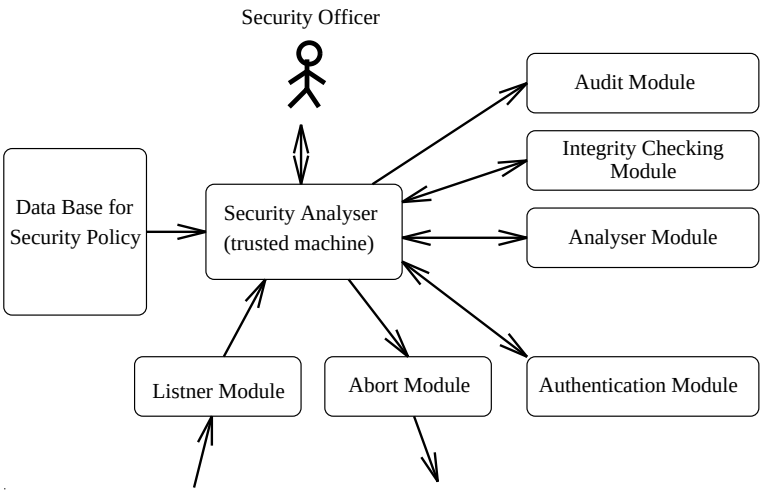

Figure 3: Architecture of the NSP

modules for checking specified security violations (such as integrity, authenticity, etc.). The abort module closes the illegal connections.

The audit module not only records all detected violations of security but also can be used as a data base for an intrusion detection system.

5 Intrusion Detection Systems

Audit is often used to monitor usage of computer and communication resources within a single security domain (single LAN). Anderson [1] argued that illegal activity (also called intrusion) can be detected by the analysis of audit records. Anderson also identified three classes of users: legitimate users, masqueraders, and clandestine users.

A legitimate user is a person who is authorized to access and use system and network resources. Legitimate users who are involved in an illegal activity by misusing their privileges are called *misfeasors*. *Masqueraders* are illegal users who subvert a computer access control system and exploit authorized users' accounts. *Clandestine* users try to take over control over protection mechanisms (for example, by suppressing the audit).

An Intrusion Detection System (IDS) looks through the audit data for a possible evidence of intrusion. So the collection of audit data is crucial to the IDS. The more thorough audit data the better are IDS decisions about real or imaginary intrusions. On the other hand, the more audit data the more computing resources (time and storage) are consumed by the IDS analysis. Automated IDS analysis of audit data provides a practical solution for an efficient processing large amount of audit data. Automated IDS analysis of audit data allows not only to identify security violations but also to tract illegal users in real-time.

While designing an IDS, it is possible to apply many different intrusion detection tools. There are three main classes of intrusion detection tools: statistical anomaly detection, rule-based anomaly detection, and rule-based penetration identification.

Statistical anomaly detection. These tools are based on statistical analysis of measurements (kept as the audit trail) of a user activity. If an intrusion detection is based on a single characteristics (event), we deal with *threshold detection*. Applications of the threshold detection can be found in NADIR [6] and MIDAS [13]. The alternative approach is to use *complex statistical profiles* to identify expected behaviour of users or user groups.

Rule-based anomaly detection. The statistical anomaly detection uses statistical formulas

to detect abnormal behaviour. On the other hand, the rule-based detection applies a set of rules to represent and store the usage patterns. This approach was used in Wisdom and Sense (W&S) [16] and in Time-based Inductive machine (TIM) [2]. W&S was developed at Los Alamos National Laboratories and can monitor the operating system as well as the application layer. The detection process can be done in real-time or off-line (in batch mode). Rules can be generated automatically from the past audit data or entered manually. TIM focuses on patterns of event sequences. Like W&S, it automatically generates rules by the analysis of the past audit data.

Rule-based penetration identification. This detection system is an expert system whose rules fire when the audit records indicate a suspicious (if not illegal) user activity. The rules may recognize a single penetration event or a sequence of events that represents an entire penetration scenario. A good example of such a system is IDES [8] developed by SRI International. IDES includes a rule-based component which allows to represent suspicious behaviour based on the site-specific security policy, known security flaws and past (detected) intrusions.

6 Secure Communication via (Insecure) Internet

As the value of information continually increases, more and more people attempt to illegally obtain the information. Furthermore, information in transit can be denied, delayed, modified or someone may masquerade their identity to illegally view information.

Authentication is a countermeasure against active attacks. An important way to classify authentication is by the number of participants. This method of classification yields: one-way, two-way and m-way authentication. One-way of authentication is where one entity is authenticated by another, e.g., user login and authentication of a workstation at a boot up stage. Two-way authentication is where two parties mutually authenticate each other, e.g., user-user, user-server, and server-server authentication. M-way authentication, also called conference authentication, is where m entities authenticate each other, e.g., processes carried out for synchronous remote meetings and concurrent group work.

This section describes common cryptographic techniques that can be applied to establish a secure communication channel in the network.

6.1 Key Distribution Between Two Parties

Most of the cryptographic techniques can be applied only if there is a secret and pre-arranged pair of cryptographic keys which are shared between two parties. The key distribution problem is even more complex in the Internet environment where there is a large number of active parties who may wish to communicate to each other (note that if the network has N parties, there are $\frac{N(N-1)}{2}$ possible pairs). The solution to this seemingly "unsolvable" problem was given by Diffie and Hellman in 1976 [3]. Their solution (also called the DH public-key distribution system or simply the DH system) is illustrated in Figure 4. The key distribution process starts from generation of a large enough modulus N (N has to be prime) and random selection of a primitive element g (g is primitive in Z_N if $\{g, g^2, \ldots, g^{N-1}\} = Z_N - \{0\}$). The pair (N, g) is made public so Alice and Bob know it. The parties A and B choose randomly their secret keys e and d, respectively. Alice and Bob calculate their partial keys

$$y_A = g^e \pmod{N} \text{ and } y_B = g^d \pmod{N}$$

respectively. Next they exchange the calculated partial keys (y_A, y_B) via insecure channel (Internet). Now Alice computes the secret key K as

$$K = (y_B)^e \pmod{N}$$

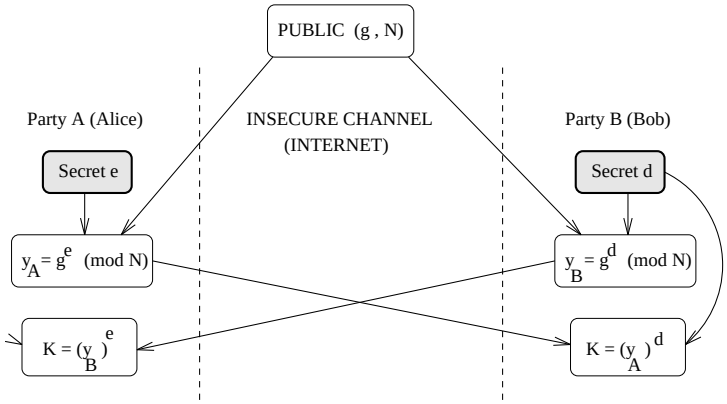

Figure 4: Information flow in the Diffie-Hellman public-key distribution system

and similarly Bob as

$$K = (y_A)^d \pmod{N}$$

Note that $(y_A)^d = (y_B)^e \pmod{N}$.

The security of the DH system relaties on the assumption that the knowledge of N, g, y_A, y_B is not sufficient to calculate secret elements e, d, K (for further details of complexity issues see [12]).

Most other methods of key distributions rely on the application of so-called master keys. These keys are distributed before hand by secure channels such as registered mail or trusted messenger. Later if there is a need to establish a common session three different logical formations: star, me

b $_i$ and H_j isusedtoencryptthesessio
w $_j$, it

ring. In all three formations a set of invitations are sent out by the conference initiator and a series of replies are returned determining the conference participants. With the star formation the initiator authenticates each participant, and each participant authenticates only the initiator.

Whilest with both the mesh and ring formations every participant authenticates every other participant. Thus, the star formation is less secure, however because it uses less messages it is faster.

However in most Internet applications, secrecy is not as crucial as authenticity. For example most of the time messages sent via e-mail are not secret but the receiver would like to be sure the the message they have received is genuine. Similarly, most of the EFT applications need to have an authentication mechanism.

References

[1] J.P. Anderson. *Computer Security Threat Monitoring and Serveillance*. James P. Anderson Co., Fort Washington, PA, 1980.

[2] K. Chen, S.C. Lu, and H.S. Teng. Adaptive real-time anomaly detection using inductively generated sequential patterns. In *Proceedings of the Fifth Intrusion Detection Workshop, SRI Int., Menlo Park, CA*, May 1990.

[3] W. Diffie and M.E. Hellman. New directions in cryptography. *IEEE Trans. Inform. Theory*, IT-22(6):644–654, November 1976.

[4] M. R. Garey and D. S. Johnson. *Computers and Intractibility A Guide to the Theory of NP-completeness*. W.H. Freeman and Company, 1979.

[5] A. Gościński. *Distributed Operating Systems the Logical Design*. Addison-Wesley Publishing Company, Inc., 1991.

[6] B. Hubbard, T. Haley, N. McAuliffe, L. Schaefer, N. Kelem, D. Walsott, R. Feiertag, and M. Schaefer. Computer system intrusion detection. Technical report, Trusted Information Systems, Inc., RADC-TR-90-413, Final Technical Report, December, 1990.

[7] K. Koyama and K. Ohta. Identy-based conference key distribution systems. In *Advances in Cryptology - CRYPTO'87, Lecture Notes in Computer Science (C. Pomerance (Ed))*. Springer-Verlag, Vol.293:175-184, 1987.

[8] T.F. Lunt. IDES: An inteligent system for detecting intruders. In *Proceedings of the Symposium on Computer Security, Threat and Countermeasures, Italy*, 1990.

[9] R. Molva and E. Rutsche. Application access control at network level. In *Proceedings of the 2nd ACM Conference on Computer and Communication Security, Fairfax, Virginia*, pages 219–228. The Association for Computing Machinery, November 1994.

[10] Government of Canada. *The Canadian Trusted Computer Product Evaluation Criteria - Version 3.0e*. Canadian System Security Center, Communications Security Establishment, Government of Canada, 1993.

[11] P. Rolin, L. Toutain, and S. Gombault. Network security probe. In *Proceedings of the 2nd ACM Conference on Computer and Communication Security, Fairfax, Virginia*, pages 229–240. The Association for Computing Machinery, November 1994.

[12] J. Seberry and J. Pieprzyk. *Cryptography: An Introduction to Computer Security*. Prentice Hall International, 1989.

[13] M.M. Sebring, E. Shellhouse, M.E. Hanna, and R.E. Whitehurst. Expert systems in intrusion detection: A case study. In *Proceedings of the 11th national Computer Security Conference, Baltimore, MD*, 1988.

[14] A. Shamir. Identity-based cryptosystems and signature schemes. In *Advances in Cryptology - CRYPTO 84, Lecture Notes in Computer Science (G.R Blakley and D. Chaum (Eds))*. Springer-Verlag, Vol.196:47-53, 1985.

[15] A. Tanenbaum. *Computer Networks*. Prentce-Hall International, Inc., 1981.

[16] H.S. Vaccaro and G.E. Liepins. Detection of anomalous computer session activity. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 280–289. IEEE, May 1989.