Authorisation and Privacy in a Networked World

Lawrie BROWN*
Department of Computer Science,
UC UNSW, Australian Defence Force Academy,
Canberra ACT 2600 Australia

Abstract

There is a growing range of tools and techniques available for providing authorisation and privacy. These are of increasing importance in a world of growing internetworking which can span many administrative domains. However for a variety of reasons, technical, commercial, legal and political, these tools are nowhere near as widespread as they could be. This talk will survey some of the available tools and techniques. There are simple extensions to common utilities such as telnet and ftp for closed user group use; the more general IETF work defining a Common Authentication Technology; the Kerberos authorisation system which provides an organisation wide user authorisation system; and finally the competing secure email technologies PEM and PGP. I will attempt to place these tools in context, and indicate their availability and ease of use.

1 Introduction

In this talk I will discuss: why networks are insecure, issues involved with providing appropriate user authentication, and the provision of privacy, particularly of email.

2 Network (In)Security

Traditionally it has been assumed that computer networks are secure. With the increasing growth of internetworking anyone with packet monitoring equipment can snoop on network traffic (Fig 1). They could monitor the logins and passwords on a local network segment. This equipment can be as simple as a PC equipped with an ethernet card, and some software to run it in promiscuous mode to accept all traffic. A recent (Feb 94) example of this occurring in the US has raised concern right up to congressional level. It involved penetrating a number of systems using a known security hole, and then running monitoring programs to capture the first 128 characters of ftp, tftp, telnet and rlogin sessions, which usually included login names and passwords. The attack employed a number of measures to evade detection, and some thousands of systems were believed compromised [1] [2]. Another well-known area of insecurity involves the use of X-terminals, which are susceptible to rogue clients which can take a copy of the current screen, or capture all key-strokes, unknown to the user. The degree of vulnerability depends on the authentication scheme being used, with the default xhost mechanism providing no protection against other users

*Email: Lawrie.Brown@adfa.oz.au

Figure 1: Network (In)Security

on authorised systems [3]. With the advent of Information Superhighways, it is obviously essential that these sort of problems should not occur, or it will adversely impact the acceptance of this technology [4]. It is also not merely sufficient to legislate against a practise that is easy, one only need consider the problems of monitoring mobile phone conversations as a case in point. Technical solutions are required that provide known levels of security.

3 User Authentication

The first task of any computer security system is the need to validate the identity of the user wishing to use the facilities. Traditionally this has been done using a username and password. Since local terminal connections are assumed to be under the control of the same organisation as the system, some level of security can be expected. However as soon as remote network connections are considered, the information may be transiting a number of networks, administered by a number of organisations. At this point the problem of monitoring can become significant. The basic approach to overcoming the problem of network eavesdropping is to ensure that the same authentication information is never used twice. Some alternatives when a small number of hosts are involved include: one-time passwords either on a paper-list, or token generated; or a challenge-response scheme with pre-arranged passwords. When security over a distributed network of servers and clients is needed, a trusted authentication key server is required.

3.1 One-time Passwords

One-time passwords can be used for authentication with pre-arranged systems. The passwords may be generated by some one-way or random algorithm and distributed in printed form (as with home banking systems), or there may be a computer or token used to generate the passwords. The best known of these for remote computer logins is the S/KEY system [5]. This uses a one-way function f to generate the passwords used. Given a secret password s, and number N, the initial password stored on the system is:

$$p_0 = f^N(s) \tag{1}$$

The first password used for remote access is:

$$p_1 = f^{N-1}(s) \tag{2}$$

On each subsequent access, the previous password in the series is used.

$$p_i = f^{N-i}(s) \tag{3}$$

Figure 2: Third Party Authentication Scheme

The system applies f once to p_i to check against the stored password p_{i-1}. Following successful access p_i is stored for checking the next login. The user may either have a pre-printed list of passwords, or use a small program on a PC to compute the required password. Since f is a one-way function, knowing p_i doesn't help an attacker determine p_{i+1}. The function f used in S/KEY is MD4 [6] [7]. Obviously when i reaches 1, the password must be reset. Hence this scheme is best used for occasional access.

Other variations on this theme use some form of smartcard or token to either supply a time-varying password, or to compute a suitable password or challenge-response value.

3.2 Challenge-Response Authentication

Challenge-Response is another basic technique used to ensure a password is never sent in the clear [8]. Given that a client and a server share a key k, the server sends a challenge vector x, the client encrypts it with the private key and returns $E_k(x)$, and the server verifies the response using a copy of the private key. In its simplest form, keys are established before secure communication is required. This technique has been used as the basis for security extensions to telnet, ftp, and TCP/IP, and are appropriate for use in closed-user groups with a small number of systems [9] [10]. In a distributed network, challenge-response is used as a component of the authentication protocol, with keys stored in a central trusted key server.

3.3 Authentication Key Servers

An Authentication Key Server uses a key server trusted by all clients, where each client shares a key with the server. The server can authenticate a client, who can subsequently request a ticket authenticating it to another server client in order to obtain a service. This ticket is sent to the other client as proof of its identity and right to the service requested (Fig 2) A number of schemes have been proposed, based on either private-key, or public-key approaches, including schemes such as Kerberos and SPX among others [11].

3.4 Kerberos

Kerberos is a trusted key server system developed by MIT, which provides centralised third-party authentication in a distributed network [12]. It is probably the best known Authentication Key Server scheme. Using it provides access control for each computing resource, in either a local or remote network (realm).

Figure 3: Kerberos Initial User Login

Figure 4: Kerberos Remote Service Request

A Key Distribution Centre (KDC) which contains a database of principles (customers and services) and encryption keys, is used to provide non-corruptible authentication credentials (tickets or tokens) to its clients. There are two phases to the authentication process:

Kerberos - Initial User Login When a user first logs on to a workstation an initial ticket is requested from the KDC which is used as basis for all remote access requests (Fig 3).

Kerberos - Remote Service Request Subsequently when a user requests access to a remote service, a ticket is obtained from the KDC protected with the remote key validating the user to that server. It is sent with the request to the remote server. All of the requests and tickets flowing between clients, and the KDC can be encrypted to prevent monitoring or modification (Fig 4).

Currently there are two versions of Kerberos available. Kerberos v4 is restricted to a single realm, but has been in production use for some time. Kerberos v5 allows inter-realm authentication, and is in early release, but not yet common. Kerberos v5 is an Internet standard, specified in RFC1510, and used by a growing range of utilities. To use Kerberos, you need to have a KDC on your network, and you need to have Kerberised applications running on all participating systems. A major problem to using Kerberos here is caused by US export restrictions which restrict source distributions outside the US. Either a binary distribution, for authentication purposes only with obscured entry points, must be obtained; or the crypto libraries must be reimplemented locally for the source distribution.

3.5 IETF Common Authentication Technology (CAT)

Given the range of authentication technologies available, there was a need for a standard way of invoking security services by applications. The IETF formed the CAT working group to provide an interface specification for accessing generic security services. Initially its designed to use Kerberos v5, though other mechanisms can be supported. It is now an Internet standard, specified in RFC1507, RFC1508, RFC1509, RFC1511. It is not widely used yet, but will be in future.

4 Cryptographic Algorithms

An essential component of many of the techniques being discussed is the use of some cryptographic algorithms. There are a number of classes of algorithms:

Private-key (symmetric) encryption algorithms are used for encrypting messages using a single shared key. Since both sender and receiver have equivalent rights, these schemes are symmetric. They may be used to provide secrecy, and for message authentication. Common algorithms include: DES, FEAL, IDEA, LOKI. The specifications for these ciphers are public, though some have patent/licencing restrictions.

Public-key (asymmetric) encryption algorithms are used for encrypting or signing messages using a public encryption (signature verification) key, and a private (secret) key. Since only the receiver knows the private key used for decryption (or signature creation), these schemes are asymmetric. They may be used to provide secrecy, message authentication, and verification of message origin. The best known scheme is RSA, other schemes include ElGamal, Rabins and others, all related. RSA is patented in US/Canada by RSA Data Security Inc., who currently are licencing its use in US/Canada. Its legal status in rest of the world is unclear, as it may or may not be covered by patent law. Also US export restriction hinder the distribution of encryption technology.

Digital Signature algorithms are used to sign a message to verify its correct delivery, and possibly (if an asymmetric algorithm is used) guarantee non-repudiation of origin. Examples include: RSA, ElGamal, DSA. Recently RSA Data Security Inc. have announced the availability of RIPEM-SIG, a US export approved signature only version of their RIPEM reference library. Although available in binary only (to prevent evil foreigners trying to use encryption capabilities), it does at least greatly ease the problem of providing authentication services in a uniform manner, world-wide [13].

Hashing algorithms produce a fixed size digest or summary of a message, which is then signed to form a digital signature. Examples include MD2/MD4/MD5, HAVAL, SHS.

A comprehensive overview of algorithms in all these classes and others may be found in Schneier [14], as well as in other texts [15] [8].

5 Privacy

Once users are authenticated, the other primary concern is the need to protect the contents of messages by either physically preventing access to the communications link, or cryptographically scrambling the data sent. The first approach

Figure 5: Email Encapsulation

may be appropriate when a single organisation controls the entire communications infrastructure (eg Defence), but will not work in an internetworking environment. If cryptographic algorithms are used, there are a number of possible public or private key encryption algorithms available. Also a key distribution scheme will be required to determine how the keys used by these algorithms are supplied to the communicating parties. The provision of privacy will be needed for a number of services, including remote access, file transfer, and email. I will concentrate on the latter here.

6 Secure Electronic Mail

Email is one of the most widely used and highly regarded network services. Currently message contents are not secure, they may be inspected either in transit, or by suitably privileged users on the destination system. That is, they have a security level about equal to a postcard. To improve this to at least the equivalent of mail sealed in an envelope, we need mechanisms to provide the following privacy enhancement services:

confidentiality to provide privacy for a message (protection from disclosure)

authentication of the sender of the message

message integrity to provide protection for the message from illegal modification

non-repudiation of origin to ensure that the sender of a message cannot deny having originated the message, provided an asymmetric digital signature algorithm is used

Email encapsulation is used to provide these features so that conventional mail delivery systems can be used without changes. Encapsulation means wrapping the original message with additional information to provide these services, and then treating this enhanced message as a mail message for the delivery system (Fig 5). Hence only the mail user agents need to be extended to obtain these services.

6.1 PEM

PEM (Privacy Enhanced Mail) is the Internet standard for security enhancements to Internet (RFC822) email. It was developed by a Working group of the IETF, and is specified in RFC1421, RFC1422, RFC1423, RFC1424. It provides:

confidentiality using DES encryption in CBC (Cipher Block Chaining) mode

integrity using a DES encrypted MIC (Message Integrity Check) generated using either MD2 or MD5

authentication using either a DES (with an on-line Authentication Key Server) or RSA encrypted MIC

non-repudiation using an RSA encrypted MIC

PEM Key Management can use either a central key server (eg Kerberos) which requires access to on-line server, or it can use public-key certificates signed by a Certification Authority (CA). The CAs form a hierarchy to permit cross-validation of certificates. X.509 Directory Service Strong Authentication [16] is used to protect key certificates. The CAs must be licenced by RSA Data Inc., who until recently have only licenced CAs in US/Canada. As the official Internet standard, this will undoubtably become the method of choice for organisations, once the current uncertainties over its use are resolved.

6.2 PGP

PGP (Pretty Good Privacy) is a widely used de facto secure email standard, developed by Phil Zimmermann. It is available on Unix, PC, Macintosh and Amiga systems, and most importantly, its free!!!! It provides:

confidentiality using the IDEA encryption algorithm [17]

integrity using an RSA encrypted MIC generated using MD5

authentication and non-repudiation using an RSA encrypted MIC

A example of a PGP signed message is given in Fig 6. It uses a grass-roots key distribution method, where trusted introducers are used to validate keys. Hence no certification authority hierarchy is required, although one could be formed if desired. All PGP functions are performed by a single program (pgp), which must be (has been) integrated into existing email/news programs. Each user has a keyring of known keys containing their own public and private keys (protected by a password), public keys given to you directly by a person, and public keys signed by trusted introducers. Their private key is used to sign/decrypt your messages, and the collection of public keys is used to validate/decrypt messages received. Possible problems with using PGP include its unlicenced use of RSA in US/Canada in contravention of patent, although there is a commercial licenced version available in US/Canada. Its use outside the US is probably legal, although the exact status is murky. This will probably remain the method of choice for groups of individuals who wish to communicate securely, but who are less concerned with global interoperability.

7 Conclusion

A need has been identified to improve the security of network services. In this talk I have discussed a range of privacy and authentication issues to address this need.

8 References

[1] CERT, "Ongoing Network Monitoring Attacks," CERT Advisory, CA-94:01, 3 Feb. 1994.

[2] ASSIST, "Network Monitoring Attacks," in *Computer Fraud & Security Bulletin*, Mar. 1994.

[3] D. Sheldrick, "Security and the X-Window System," in *UNIX World*, Jan. 1992.

[4] P. Wallich, "Wire Pirates," *Scientific American*, 270, no. 3, pp. , Mar. 1994.

[5] N. M. Haller, "The S/Key One-Time Password System," Bellcore, ISOC symposium paper, 1994, ftp://thumper.bellcore.com /pub/nmh/docs/ISOC.symp.ps.

[6] R. L. Rivest, "The MD4 Message Digest Algorithm," *Advances in Cryptology - Crypto'90*, 537, pp. , 1991.

[7] R. L. Rivest, "The MD4 Message Digest Algorithm," MIT and RSA Data Security Inc., RFC1320, Apr. 1992.

[8] D. W. Davies and W. L. Price, *Security for Computer Networks*. New York, John Wiley and Sons, 1989, (2nd edn).

[9] L. Brown, "Secure Remote Login - the SECLOG option," in *AUUG 90 Conference Proceedings*. Sydney, NSW, Australia: Australian UNIX Systems Users Group, pp. , Sept. 1990.

[10] L. Brown and M. G. II. Jaatun, "Secure File Transfer Over TCP/IP," in *Proc. IEEE Tencon-92*, Melbourne, Australia, Nov. 1992, Also available as Dept. of Computer Science, UC UNSW, Australian Defence Force Academy TR CS2/92.

[11] D. Gollman, T. Beth and F. Damm, "Authentication Services in Distributed Systems," *Computers & Security*, 12, no. 8, pp. , Dec. 1993.

[12] J. G. Steiner, C. Neuman and J. I. Schiller, "Kerberos: An Authentication Service for Open Network Systems," in *Proc. Usenix Winter Conf.*. USENIX Assoc., pp. , 1988.

[13] M. Riordan, Exportable RIPEM/SIG Available, 1994, ftp://ripem.msu.edu /pub/crypt/ripem/ripemsig/posting.

[14] B. Schneier, *Applied Cryptography - Protocols, Algorithms and Source Code in C*. New York, John Wiley & Sons, 1994.

[15] J. Seberry and J. Pieprzyk, *Cryptography: An Introduction to Computer Security*. Englewood Cliffs, NJ, Prentice Hall, 1989.

[16] G. Dickson and A. Lloyd, *Open Systems Interconnection - Computer Communications Standards and GOSIP Explained*. New York, Prentice-Hall, 1992.

[17] X. Lai, J. L. Massey and S. Murphy, "Markov Ciphers and Differential Cryptanalysis," *Advances in Cryptology - Eurocrypt'91*, 547, pp. , 1991.

```
-----BEGIN PGP SIGNED MESSAGE-----

May all your signals trap
May your references be bounded
All memory aligned
Floats to ints be rounded

Lawrie
-----BEGIN PGP SIGNATURE-----
Version: 2.3
iQBzAgUBLdl1RILpoub8ek7fAQF2nwLuJwVPh8ii
FrksXSCe6z37ZdV37pXvsYyz0WAnCBCdpu55yId5
/kVhmvusTo10zUHPssPwB99TQq9YsduSfkVeILjf
JNJEuUWQkJl8dWvaB+IIEEodF0Xpbc23krnu0A==
=hn90
-----END PGP SIGNATURE-----
```

Figure 6: PGP Signed Message