

The Internet Opens Up

Electronic Mail

By Matthew Gream

If there really is someone else reading your e-mail, is encryption the answer?

There are some fairly strong allegations being flung around about Internet mail security. Some would have us believe the CIA is scanning and logging mail, while others closer to home have suggested our own Australian intelligence organisations are carrying out similar exercises. The truth is, we really don't know what the situation is – and we're unlikely to see anyone come out and admit it either. What we can be sure of though, is that if such operations are being carried out, then we can start pointing fingers at those agencies chartered with Signals Intelligence (SigInt). I'm specifically thinking of – as are many others – the US National Security Agency (NSA). It is this particular agency which carries out SigInt in the form of surveillance of radio- and electromagnetic emissions, along with providing Communications Security (ComSec) to US Government bodies. James Bamford's expose of this organisation, "The Puzzle Palace", reflects as much the secrecy surrounding it as the "codebreaking" operations forming an inherent part of the SigInt collection. In his book, Bamford describes how the NSA spent some of the best years of the Cold War running an extensive operation, taking in and examining every "cable" before it left the US bound for foreign soil. In addition to this - but for a shorter period

they carried out a letter-opening campaign with the co-operation of the US Postal Service. Letters destined for "red" countries were subject to scrutiny by rooms of letter openers. While these particular activities have long since stopped, in recent times, certain comments made by NSA have shown they know what many "popular" personalities are up to on cryptologic and digital politics Usenet newsgroups.

Fair enough! Usenet is a public broadcast medium with millions of participants, so it's perhaps not surprising they should take an interest. Given the increasing population of the Internet, they may be taking even more of an interest.

What makes surveillance on the Internet intrinsically different from letter opening, cable examination or telephone call monitoring is that it's so easy. Gathering the raw information requires little more than an average Unix workstation plugged into a network of interest.

The Medium

Ethernet, the almost universally utilised networking data link layer, is a broadcast medium, so every frame placed on a network is readable by all others on the same network. Although certain wiring topologies can alleviate this problem, they are usually employed at the lower end of the hierarchy. Having a medium supporting passive surveillance means half the battle is won.

Putting these abstract Ethernet frames into some order proves just as easy. In fact, software specifically constructed to reassemble TCP streams has featured in the most

recent spate of wide scale Internet attacks. The item in question is geared to log the first few kilobytes of each login session, which usually contains a "login" and "password" pair for the remote host. To gain the widest possible audience, a key component of these attacks has been to gain access to hosts that live on backbones and other high volume, important transit networks. Subsequently, single attacks have resulted in the compromise of hundreds, even thousands, of local and remote hosts. The software was written by no more than a wily hacker, and derivatives currently in circulation grab and retrieve electronic mail containing "interesting" keywords. While a hacker's favourite list may start with "password", "login", "security" and "secret", a government counter-terrorist agency maybe more interested in those with "sedition", "terrorism", "explosive", "black bag" and so on.

If this is what (largely) unskilled hackers with minimal resources can achieve, then the potential for well-funded government agencies is well and truly open to speculation. Unfortunately government agencies aren't the only ones likely to have an interest in reading the electronic mail of others. As in other areas of society, corporate bodies and dubious individuals tend to create the greatest threat to privacy and security. Part of the reason we see this being amplified on the Internet is that of a topology by which an item of electronic mail may pass through any number of networks.

It should not be surprising that somewhere along the chain someone will not subscribe to your set of values when it comes to privacy. Your

own system administrator may well be on the lookout for good lunch room gossip material, and a mail spool represents a likely target.

The ease with which electronic mail can be passively monitored is nothing compared to the ease of active forgery. After all, the native Internet mail transport agent is SMTP (Simple Mail Transport Protocol) and contains virtually no authentication mechanisms. There is only one field a local SMTP agent inserts into a mail header that you can be sure of, but the uninitiated wouldn't pick the difference between a fake and the real thing.

Primary e-mail environment

If you're crying foul over my focus on the Internet, then you'll be pleased to know the Internet is but only one fine example because of its size and use. It also happens to be the primary electronic mail environment. The problems previously discussed are just as applicable to self-contained internal networks and differing mail agents to some degree. For example, one particular Australian Government Department has recently had problems with an employee who was able to gain unauthorised access to electronic mail. Quite simply, the situation is that most users, unless specifically informed, are just not aware of these problems. In the physical world, there is the concept of postal mail travelling through the hands of real people. We can readily see that an "open" postcard is a free-for-all read, and few people would consider sending anything more than a greeting in something so visible.

The technology to alleviate the problems of electronic mail security isn't exactly lacking, but could be best described as being in an early evolutionary phase and best characterised by Pretty Good Privacy (PGP). An item of software, originally written by Phil Zimmermann, PGP has been rapidly and extensively deployed on the Internet, partly because of a need but also because of its source code availability – essential for removing doubts about the product's quality.

PGP uses Public Key Cryptography System (PKCS) and a technique known as RSA (Rivest-Shamir-Adleman, after it's authors), to provide digital confidentially and authentication.

To explain the operation of a PKCS, it's necessary to remember that conventional cryptography - the encoding and decoding of information - is employed with a single secret key known to both sender and receiver. This key must be previously transmitted in secret before it can be used to encode information. However in a PKCS, two keys exist – one being a widely published public key and the other being a private key, which is kept secret.

Certain mathematical relationships exist between the two keys. The most important is that given a message encoded by one key, it is not possible to derive what that key, or the other, is.

Sending a private message requires obtaining the recipients public key and using it to encode the message. A digital signature can also be added by encoding the message with the senders private key.

Only the recipient can decode the message because only it holds the matching private key. In verifying the digital signature, the recipient decodes the message with the sender's public key and, as no one else is assumed to know the senders private key, the sender can be reasonably sure of where the message originated. Encoding for a particular recipient isn't a necessary prerequisite to a digital signature. They can be applied just as easily to a clear message that anyone can read. This technique is frequently used on Usenet messages.

Web of trust

Digital signatures on public keys themselves are just as important as those on messages. A key signed by someone you trust, or an organisation, provides you with some assurance that the key you are using is in actual fact owned by who it says it is. This model used in PGP is referred to as a "web of trust". It differs from

other key management schemes which require signatures by upper authorities in hierarchical trees.

Following this trend of informality, PGP public key servers have sprung up across the Internet and are mostly run on a voluntary basis. Publishing a key via such a distribution service means it is available for someone to use in a private authenticated conversation without first having to obtain the key from a desired recipient. PGP is not an official standard, and by no means the only Privacy Enhanced Mail (PEM) software available. In typical Internet fashion though, it has risen to de facto standard and is enjoying phenomenal use. The PEM that is standardised and has software support is, in contrast, floundering. The real problem to overcome, regardless of which type of PEM, is the need to integrate and automate encryption, signature and key management functions in popular mailers and then deploy this software on a wide scale. Although ad hoc scripts and patches have been developed to aid in the use of PGP with popular mailers, integration is problematic and far from seamless.

Commercially, Apple has made a start by offering an RSA implementation "built in" to it's operating system. Unfortunately, US export restrictions on cryptographic software (it's defined as a munition with potential for military application) mean the maximum permissible key sizes are restricted. Key sizes are a measure of security, and Apple's 512 bits falls below the 1024 bits many recommend. •

Matthew Gream

<M.Gream@uts.edu.au> is a computer systems engineer specialising in the communications and information security domain. For a list of Frequently Asked Questions (FAQ) on where to obtain PGP and how to use PGP, send electronic mail to the author.