Keeping e-mail secure
Some useful tips for keeping snoopers out and sensitive data in

Stephen Satchell
Developer and author
October 2000

The hot debate surrounding the discovery of the U.S. FBI's Carnivore system has one beneficial effect: Companies have been given a huge heads-up that e-mail can be snooped without warning, notice, or even detection. What the FBI can do, so can a "black hat" performing industrial espionage. This article addresses common-sense approaches for preventing your electronic mail from compromising your company or your business. In particular, it discusses the new addition to the mail protocols that enhances security of electronic mail.

Contents:

Background

Who could snoop your electronic mail? Anyone.

Law enforcement agencies use the Carnivore program to implement a wiretap order on Internet-based electronic mail of a suspect in an investigation. Carnivore does this by "snooping" the network on which a mail server is connected, looking for packets that meet a pre-programmed set of filter criteria. Packets that meet the criteria -- criteria that presumably limit the information captured to that specified in a Court order -- are then stored onto Exabyte tape for later processing. When installed properly into an ISP site, Carnivore is completely undetectable.

There is no rocket science in the construction of Carnivore. The core of the thing is a commercially available network packet monitor/analyzer -- or "sniffer" -- a network diagnostic tool that's used by network administrators everywhere. During Congressional hearings about Carnivore, network administrators testifying to Congress demonstrated just how easy it is to tap into electronic mail using these tools without disturbing servers or networks -- without setting off alarms of any kind. A number of vendors sell powerful sniffers, but anyone can use Open BSD or Linux for free, along with readily available programs and libraries, to do the same job just for the price of an entry-level personal computer and the effort of downloading the source and compiling the packages. The U.S. FBI is just a new member of a large group abusing diagnostic tools for non-diagnostic purposes (such as information espionage).

An espionage agent using a packet sniffer to snoop Internet activity does so with a specific purpose in mind. Like any data mining activity, the "take" has to be managed with that purpose in mind, lest the

poor agent succumb to being buried in the stuff that has been snooped or starved of the very information they were looking for.

The snooping methods can be broken down into three classes. The most voluminous method is *full-content snooping*, where the full text of messages is intercepted and stored. The next level up in abstraction is *detailed traffic monitoring*, where the sending and receiving addresses are captured, and perhaps some characteristics of the message (length, message type, or other easily detected attributes) are stored. The most basic level of monitoring is *site-to-site* (or *site-level*) traffic monitoring, where the traffic volume between IP addresses is monitored.

Full content snooping

Reliable, full-text capture by a third party is very difficult, technically and analytically, and useful only when the mail contains secrets. The technical issue is whether the sniffer can keep up with the data stream. The analytical issue is that someone has to go through the "take" and pick out the nuggets of useful information.

Full-text capture is also the easiest snooping to block: Just encrypt the contents of electronic mail with any number of packages. Many people have been using Phil Zimmerman's PGP (PrettyGoodPrivacy), currently available for free download at Network Associates, Inc.'s Web site (see Resources). Commercial versions are also available at NAI. Other vendors have made encryption packages available. The S/MIME standards are described in RFCs 2631 through 2634 (version 3) and an earlier effort in RFCs 2311 and 2312. (See Resources.)

Because there are several incompatible encryption products available, you should plan your encryption implementation carefully. For example, you should be sure that all parties share a common encryption algorithm so that a message encrypted on one system can be decrypted on the others. Get it to work *before* you need it.

The next question is how strong an encryption you should use. The goal of encryption is not to absolutely protect data; rather, it is to increase the time and cost of "cracking the code" to the point that the return on the effort is too small to be worth expending the effort. For public-key cryptography, a 512-bit key is considered less effective, while a 2,048-bit key is considered to have a significant life.

Detailed traffic monitoring

Intelligence doesn't have to be detailed to be crippling. For example, journalists can tell when something is up in a particular government agency in Washington, D.C. by the number of lights burning in the office buildings at night. In the case of signals intelligence, just the *number* and the *size* of signals between points A and B can suggest not only that something will be happening, but the nature of that something as well.

Law enforcement puts its faith in pen-register orders and trap-and-trace orders. The legal barrier to getting a court order for a full ("Title III") wiretap is a high one in the courts: The law enforcement officers have access to *everything* said on the telephone line, so the requesting authority needs to show an overriding need to invade. The barrier for pen-register and trap-and-trace orders is much lower because the information gleaned from such orders is not considered by the courts to be "content," but rather "addressing information." The courts have already extended this concept to Internet transmissions. (Whether the implementation of Carnivore captures content beyond that allowed by Court order -- as a number of privacy advocacy groups claim -- is the core of the current debate, and is beyond the scope of this article.)

Encryption of the body of the message can't protect the addressing information: The protocols used for transferring mail separate the address from the body, much as the outside of a mail envelope exposes the address information while the content is tucked inside. There's even a question whether capturing address information is considered "wiretapping." In the United States, courts have made a distinction between content capture and address ("digit capture" in the case of the telephone) in terms of the type of search warrant required. Because of this bifurcation, you (or your company) may find yourself hampered in prosecuting an espionage agent who clearly just "reads addresses" and not content.

Content can be important to a industrial espionage agent only insofar as it classifies whether the message is plain text or cipher text. If a company is in the habit of encrypting "delicate" messages, the agent can measure separately the number, sender, and recipient of such messages and derive useful information from that. One way to reduce the effectiveness of such classification is to encourage widespread use of encryption, even for less sensitive material. This robs the espionage agent of even the most crude message sorting method.

Site-level traffic monitoring

This is the simplest form of monitoring, measuring the amount of traffic between IP addresses from a tap at a convenient point, as one might monitor the service entrance of a building or telephone equipment room. Off-the-shelf hardware and software that can perform this task is readily available; agents can even download “free” software that is more than up to the task. The monitor captures only the header of IP packets, and usually counts the number or data volume between IP address pairs on the network.

Defeating this sort of monitoring is easy. For leased-line circuits, consider using an external compression and encryption box on each end of the circuit. For T1 circuits, not only does this protect your traffic, but it can increase the data rate from 1.544 Mbps to over 3 Mbps.

Encryption, encryption, encryption

Many of our suggestions involve the use of encryption in one way or another. The decreasing price of computation makes encryption more and more feasible for protecting yourself from espionage. Let’s look at the possibilities.

Virtual private networking (VPN) is the process of setting up a "tunnel" through the Internet from one point to another. The network is private because encryption is used at the two ends, to hide the information passing through the public medium (i.e., the Internet) from prying eyes. Cisco and Lucent build VPN capability into their router products, while numerous other manufacturers offer appliances. Microsoft includes VPN capability in its Windows product line. Mail sent through a VPN tunnel is effectively protected from snoopers. Other applications that use the tunnel -- such as Web browsing, file transfers, and IRC chats -- are protected as well. Indeed, when a VPN is used as a general-purpose virtual link, the VPN is transparent to all TCP/IP applications.

The tunnel concept has been extended to mail transfers generally to increase the security of *all* mail traffic. RFC 2487 describes a method of extending the Simple Mail Transfer Protocol (SMTP) to use Transport Layer Security (TLS -- also known as Secure Socket Layer, or SSL) tunnels. This means that mail sent from a TLS-aware mail client to a mail server, or between mail servers, is encrypted by default -- addresses and content. The mail client doesn’t have to understand how to decrypt individual mail messages in such a system. Further, the addressing information is hidden from eavesdroppers.

How does it work? The handshake between mail transfer agents, the software that actually exchanges

mail from server to server or from the client to the first server, includes a special indicator that says, in essence, “I know how to do encrypted transfer.” The two pieces of software exchange this special token, build an encrypted channel, and then proceed as they always do.

The key is that *all* information -- address and content -- is transmitted in an encrypted channel. The encryption is as strong as the two servers allow, which means that all mail ends up being encrypted between two TLS-capable mail transfer agents, and not just those messages considered delicate. The addresses are hidden as well, so that if the encryption isn't cracked in the virtual channel, the only information available to the espionage agent is site-to-site traffic volume, with no indication of what percentage of the traffic is sensitive.

The certificate exchange system used in TLS-enhanced SMTP can also protect against forged headers that cause people to leak secrets. Users can check to see if the source of the e-mail follows a path that is legitimate for that person.

The servers known to provide RFC 2487 capability are Sendmail version 8.11 and Microsoft Exchange Server 5.5. Qmail by Dan Bernstein augmented with a patch by Frederik Vermeulen is an experimental implementation.

TLS-enhanced IMAP, POP3, and ACAP

If your company has mail users outside of the building, then you also need to look into RFC 2595, which extends the TLS capability to communications between mail servers and mail clients. As with the server-to-server communications protection, the server-to-client protection ensures that, where possible, mail is not subject to any form of interception or significant traffic volume measurement.

Today, there are no commercial mail clients for Windows that implement RFC 2595. Expect to see new products by the end of the year from Qualcomm/Eudora and from Rit Labs, just to name two.

Security is more than just eavesdropping

This article looks at direct means for monitoring communications. But there have been Trojan programs that have infiltrated computer systems that go directly to the hard disk and send information, be it in mailboxes or in files. Some Trojans “phone home” for instructions when they penetrate a system, giving an espionage agent unauthorized access to your system, and also to your servers. That’s the focus of a different article.

Encryption techniques are only one way of protecting your data. You also have to pay attention to physical security, access control, and document storage to ensure that no espionage agent can gain access to your data.

Conclusion

Electronic mail is quickly growing into the primary business communications tool. The information that is exchanged can be sensitive: contracts, reports, medical histories, designs, even bank statements. This article shows how passive monitoring of mail can be thwarted by good habits and implementation of common-sense cures for the snoops.

The future? TLS-enhanced mail can only grow now that the bellweather mail transfer agent, Sendmail, has the capability built-in. Other mail transfer agent programs will need to be enhanced in the same way. That further reduces the possibility of casual eavesdropping as time goes on.

If you are breaking the law, don’t think for a moment that these techniques will hide your actions. Law

enforcement has other ways of catching and convicting you.

Resources

- Visit Network Associates' Web site.
- RFC 2487 describes a method of extending the SMTP to use TLS tunnels.
- RFC 2595, "Using TLS with IMAP, POP3 and ACAP," extends the TLS capability to communications between mail servers and mail clients.
- RFC 2631: "Diffie-Hellman Key Agreement Method"
- RFC 2632: "S/MIME Version 3 Certificate Handling"
- RFC 2633: "S/MIME Version 3 Message Specification"
- RFC 2634: "Enhanced Security Services for S/MIME"
- The Sendmail Consortium maintains Sendmail.org, a resource offering freeware versions of the software.
- Qmail, by Dan Bernstein augmented with a patch by Frederik Vermeulen, is an experimental implementation designed to replace Sendmail.
- Eudora is expected to come up with a commercial mail client for Windows that implements RFC 2595.
- RIT Labs offers The Bat, an e-mail client designed for a wide range of users.
- Microsoft Exchange Server also provides RFC 2487 capability.

About the author

Stephen Satchell is a long-time software developer, starting with his work on ARPAnet in 1971.He, along with co-author HBJ Clifford, wrote *Linux IP Stacks Commentary*, which describes the TCP/IP implementation contained in the Linux operating system. An early adopter of PGP, he has been working with information security for more than a decade.