The Difficulties of Tracing Spam Email

Dan Boneh
dabo@cs.stanford.edu
Department of Computer Science
Stanford University*

September 9, 2004

1 Introduction

At the FTC staff's request, I prepared this report to discuss the specific difficulties of identifying a spammer by following the electronic trail embedded in spam email.[1] This report will not discuss any investigative technique other than electronic tracing — it will not, for example, discuss tracking spammers based on insider information, spam body analysis, or other investigative techniques that do not involve electronic tracing.

This report is organized as follows. Section 2 describes how the email system works and how it is used by spammers. I survey a number of identity concealment techniques

*Affiliation is given for identification purposes only. This report is not sponsored or endorsed by Stanford University.

[1]For the purpose of this report, the term "spam" refers to unsolicited commercial email that violates the requirements of the CAN-SPAM Act. A "spammer" is a person who faces liability under the CAN-SPAM Act for the impermissible spam.

commonly used by spammers. Section 3 discusses the feasibility of successfully tracing the origin of spam email by following the electronic trail. The main question addressed is whether

cannot ordinarily be forged, but the information it provides is limited, as discussed below.

The Received header. Suppose Alice sends email to Bob. Typically, Alice's email server connects to Bob's email server using the SMTP protocol and forwards the email. More specifically, recipient Bob's email server works as follows:

1. Bob's email server waits for connections to its email port (port 25). Typically, any machine on the Internet can initiate a connection to the server's email port.
2. When a connection is established from some remote machine on the Internet, Bob's email server requests that the remote machine provide the following information: (i) the recipient's email address,[3] (ii) the sender's email address, and (iii) the email message's content (including headers and body). However, there is no requirement that the remote machine provide the sender's *true* email address; the sender's email address is easy to forge. Spammers frequently provide false and misleading information in place of the sender's true email address. In addition, spammers can insert other arbitrary false or misleading headers in the email.
3. Once Bob's email server receives the email message, it adds the message to Bob's inbox. However, before doing so, it inserts an important header at the start of the email called the Received header. This header becomes the first Received header in the email and it identifies the network address[4] of the remote machine that initiated the connection (potentially, the network address of Alice's email server).[5] Because this header is inserted by Bob's email server, Bob can trust that the information this header provides is accurate.

[3] An email address is an alphanumeric string such as `'johndoe@recipient.com'`.

[4] The network address of a machine on the Internet (also known as its IP address) is a group of four numbers of the form `'192.168.1.1'`.

[5] Note that the above model assumes that the email message is transmitted directly from sender Alice's email server to recipient Bob's email server. In actuality, the information frequently passes through a series of computers after leaving Alice's email server and before reaching Bob's email server. The first Received header indicates that last computer through which the message has passed before being delivered to Bob's email server.

To reiterate, note that the first

the Phatbot network is large and designed to spread rapidly,[12] the Bobax network is much smaller and only assimilates drones that have high-speed Internet connections.[13] Phatbot drones can be instructed to "harvest" or collect email addresses to be spammed. Of crucial importance, drones in both networks can be used to send spam email. Each Phatbot drone can be instructed to start a new open proxy server on the drone machine, which can be exploited by a spammer as described in the discussion of open proxies below.[14] Each Bobax drone can be instructed to act as an anonymous open mail relay, sending a copy of a particular

using an open proxy outside the country. Unfortunately, open proxies also help spammers send untraceable bulk email.

When a spammer sends email through an open proxy, the email is forwarded from the proxy to the spam recipient. From the email recipient's point of view, the email is coming from the proxy, not the spammer's computer. The spammer's network address is nowhere in the email. Using open proxies to hide a spammer's network address is so effective that current estimates suggest that over 60% of all spam is sent through open proxies.[15] Bulk email tools such as Send-safe are capable of searching the Internet for these open proxies.

The number of open proxies on the Internet is significant. Internet services that sell lists of open proxies are currently advertising over one thousand functioning open proxies. Attachment A describes in more detail how open proxies could be used for sending spam email.

3. Open mail relays. Email messages are hardly ever sent directly from the sender's email server to the recipient's email server. Instead, email messages pass through a number of gateways called "mail relays." Each time an email message passes through a mail relay, the relay inserts a Received header at the front of the message that shows the address of the computer that connected to the mail relay. By the time the email message reaches its recipient, it contains a number of Received headers: one for every relay server through which the email message has passed. When not manipulated, the list of Received headers identifies the entire chain of relays that processed the email message. If no spammer manipulation takes place, the network address of the sender's email server is contained in the last Received header.

When a mail relay is configured properly, only certain computers may connect with it to exchange information. Such a mail relay either *accepts* email from computers on a limited list or *transmits* email to computers on a limited list.

[15] See Wood report, reference [8]. This includes open proxies on drone machines.

An open mail relay is a misconfigured mail relay that exchanges information with *any* other computer on the Internet – it accepts email from any computer on the Internet and forwards email to any other computer on the Internet. Despite an ongoing effort to shut down all open mail relays, many still exist.[16]

Spammers can use open mail relays to conceal their identities. Instead of sending email to the recipient directly, the spammer sends email to the open relay and the open relay forwards the email to the recipient. The email messages travels from the spammer to the mail relay and then to the recipient. From the recipient's point of view, the email appears to come from the open relay, not the spammer. More precisely, only the mail relay communicates with the recipient's email server. The spammer's email server does not.

Open relays do not conceal the spammer's identity as well as bot-networks or open proxies. Due to the nature of mail relay, the network address of the spammer's email server appears in one of the Received headers in the email. Nevertheless, most bulk email tools add fake Received headers to email so that the recipient cannot tell which of the Received headers in the email message contains the network address of the spammer's email server.

Note that a spammer who uses an open relay can nevertheless protect her anonymity. Spammers can obtain both the anonymity provided by an open proxy and the economy provided by an open relay[17] by routing spam emails first through an open proxy and then through an open relay. Through this route, a spam email message travels from the spammer to the email proxy, then to the mail relay, and finally to the specified recipients. Headers in the final email contain the open proxy's network address, but reveal nothing about the

[16] For example, the `openrelaycheck.com` service currently lists about 288 functioning open mail relays that were tested within the last month.

[17] There is a strong economic incentive for spammers to use open relays rather than open proxies: an open relay allows a spammer to send many more emails with the same amount of time and effort. For example, when a mail relay receives a single email message with 50 email addresses in the Blind-Carbon-Copy ("BCC") field, the mail relay sends 50 emails, one to each BCC recipient. Hence, by sending only one email message, the spammer is able to spam 50 recipients. It is the open relay that copies the message 50 times, not the spammer. Such a spammer need not maintain an expensive high speed Internet connection to send lots of emails. The spammer amplifies his connection speed through the mail relay.

spammer's identity.

4. Untraceable Internet connections. There are several ways to access the Internet through a network address that cannot be linked to an individual or a physical location. Users who connect to the Internet through public Internet cafes, through free (or stolen) wireless connections, or through certain universities' on-campus networks need not identify themselves and can therefore send messages anonymously on the Internet. Spammers may also purchase ISP roaming access using false names and untraceable payment methods. There is no way to associate such network addresses with the spammers who use them.[18]

When using an untraceable Internet connection the spammer need not hide her network address and can send email directly to spam recipients (by directly connecting to the email port on the recipient's email server). In some cases this approach must be combined with open proxies to bypass certain ISP restrictions. Using an untraceable Internet account in conjunction with an email proxy has the additional benefit of further obscuring the spammer's network address.

For further discussion of methods that spammers use to conceal their identities, please see Chapter III of the FTC's Report regarding the feasibility of a Do Not Email Registry [3].

2.2 Spamware: Bulk Email Tools

Spammers and legitimate email marketers can use software designed for sending email in bulk. There are dozens of such software tools on the market, typically costing on the order of $500. Most are designed to generate and send about 500,000 email messages per

Most bulk email tools are designed to permit spammers to hide their identities through methods discussed in the previous section: spoofing, open proxies, open mail relays, bot-networks, and untraceable accounts. It is up to an email marketer to choose which methods to use.

cybersleuths.

There are two primary approaches that cybersleuths can use: analyzing email headers and constructing computer honeypots to attract and trap spammers. I discuss these approaches and their limitations in turn below.

3.1 Header Analysis

The most straightforward approach to tracing an email message's electronic path might be to analyze its headers. After all, email headers are designed to identify the route by which the email has traveled. However, when spammers use the various identity concealment techniques discussed in the previous section, analyzing the email's header information generally will not successfully identify the spammer. This is because the spammer's email header will contain little or no information identifying the spammer's network address, unless the spammer is incompetent.

Based on my own informal study and other sources,[20] I estimate that at least 71% and, more likely, as much as 90% of spam email does not contain the spammer's network address, and therefore contains no information that identifies the spammer.

Of the four identity-concealing techniques discussed above, two of the techniques (open proxies and bot-networks) can successfully conceal the spammer's network address, and a third technique (use of untraceable Internet connections) identifies the network address of an email account that is untraceable to the spammer. For example, when a spammer sends email through an open proxy, the network address will falsely point to an innocent person. Similarly, when a spammer sends email through a bot-network, the network address also

[20] An Oct. 2003 press release from Brightmail [1] states, among other things, that "... 90% of spam is untraceable by available methods." Similarly, Wood [8] estimates that over 60% of spam email is sent through open proxies and hence is untraceable by header analysis. My own informal study of the percentage of untraceable spam email is provided in Attachment B.

falsely points to an innocent person. Finally, when a spammer uses an untraceable Internet connection, the spammer need not hide her network address. This is because there is no way to associate such untraceable email accounts with the spammer.

Only the fourth identity-concealing technique — the use of open mail relays — might occasionally contain useful information, in the sense that the network address of the spammer's email server may appear in *one* of the Received headers in the email. As described earlier, by the time an email message reaches its recipient, it contains a number of Received headers: one for every relay server through which the email message has passed. Thus, when the email message has passed through an open mail relay, one of its Received headers will contain the network address of the open relay. However, it would be difficult for a cybersleuth to identify *which* of the network addresses appearing in the Received headers is the true network address of the spammer's computer server. For example, most bulk email tools used by spammers add fake Received headers to email so that the recipient (and any cybersleuth) cannot tell which of the Received headers in the email message contains the network address of the spammer's email server.[21] Even assuming that a cybersleuth were occasionally able to identify which of the many fake Received headers is the spammer's true network address, it is likely that over time spammers would simply stop connecting to open mail relays directly, given the alternative identity-concealing techniques available to them. In such a scenario, analyzing the header for information about the spammer's network address would be rendered completely ineffective.

3.2 Honeypot Computers

In addition to analyzing the information contained in the spam's header, another method of electronic tracing available to cybersleuths is to set up "honeypots." A "honeypot" is a computer system set up to attract and trap those who attempt to exploit other people's computer systems. Cybersleuths might be able to set up honeypots to lure spammers and expose their identities. However, as discussed below, spammers are already well aware of the existence of honeypots and are implementing methods to evade honeypots. If cybersleuths, working together closely with law enforcement, were to increasingly use honeypots to trap spammers, it is likely that spammers would simply increasingly use countermeasures to avoid honeypots. Indeed, we can expect such a "cat and mouse" pattern to continually evolve in the future.

Honeypots could be applied by cybersleuths to cases in which the spammer is either using an open proxy to conceal her identity, or using a bot-network to conceal her identity. There is no point in applying honeypots to cases where the spammer is using an untraceable Internet connection since, in such cases, even if the honeypot were to successfully trace the spammer to an untraceable Internet connection, there is no way for the cybersleuth to link or tie the untraceable Internet connection to the spammer. Honeypots apply to open mail relays in exactly the same way that they apply to open proxies. For this reason, in the following section, I will briefly describe honeypots only as they apply to (1) open proxies and (2) bot-networks.

Honeypots Applied to Open Proxies. Recall that an open proxy enables spammers to fully conceal their identities by making all email messages appear to come from the proxy. A cybersleuth could set up an open proxy honeypot and wait for spammers to start using it. This fake open proxy would record the source address of all connections to it along with all tra□c routed through it. This could potentially provide significant leads for catching the spammer.

The Proxypot Project[22] is an example of an open proxy honeypot specifically designed to catch spammers. It accepts connections from any computer on the Internet, and logs all relevant information about the connection. Most importantly, it logs the address of the computer that initiates each connection. The project also provides tools to search these log files for spam activity. Note that Proxypot actually stops short of sending spam tra□c to its destination. It only logs the fact that an attempt to send spam has occurred. By blocking spam routed through it, Proxypot ensures that it does not contribute to the prevalence of spam email.

Honeypot logs can reveal the network address of a spammer.[23] Once spammers discover the open proxy honeypot, they begin to route spam email through it. Unless the spammers take extra precautions, they cannot tell that the open proxy they are using is a honeypot. A few days later, after examining the honeypot's logs, the cybersleuth can expose the spammer's network address.

Spammers are already implementing methods to evade honeypots. Although open proxy honeypots would seem to be a powerful technique for catching spammers, there are a number of significant drawbacks to this approach. First, spammers are well aware of the existence of honeypots and are implementing counter-measures to avoid them. For example, the Send-safe tool is capable of detecting honeypots by sending a test spam email to itself. Since most honeypots block spam email routed through them, the test message will not be delivered and Send-safe will stop routing email through the open proxy honeypot.

Second, spammers can completely fool an open proxy honeypot by using proxy chains. Suppose the spammer identifies three open proxies called A,B, and C. The odds are that at most only one of them will be a honeypot. The spammer then sends email by creating a path through all three servers. The email will travel from the spammer's machine first to server

A, then to server B, then to server C, and finally to the spam recipient. Now, suppose that server C is the honeypot. It only "sees" connections from server B, not from the spammer.

A second, and more powerful, spammer technique to evade honeypot detection has arisen more recently. Spammers often now design bot-networks so that the sites with which individual drones communicate are not fixed. For example, drones in the Phatbot network receive instructions using a peer-to-peer network of drones. Because the honeypot drone in such a network only communicates with a few other drones, its view of the bot-network is local and limited, and it would not have access to the network address of the bot-network administrator. Thus, as these two spammer techniques to evade detection illustrate, we can expect this "cat and mouse" pattern to play out repeatedly as sophisticated spammers increasingly use and evolve new such methods to evade honeypot detection.

4 Risk of Fabricated Evidence

Needless to say, a reward system might motivate dishonest individuals to submit fabricated evidence about the origin of spam. It is especially easy to fabricate evidence about electronic mail due to the lack of security in the SMTP protocol. In fact, the same identity concealment techniques discussed in Section 2.1 could also be used to fabricate evidence.

- The lack of email origin authentication makes it easy to create the text of an email that appears to come from an innocent party. The headers in the email would be spoofed to make it look like spam email from the innocent party. Many such fake emails could then be submitted as evidence of spamming in order to claim a reward.
- A bot-network administrator could "sacrifice" a few drones and submit evidence that these drones are sending spam. Furthermore, the bot-network administrator could easily plant false evidence on the drone machine to make its owner look like a spammer. Bot-network administrators currently have no incentive to do so, but a reward system might change that.

5 Conclusion

This report has discussed two electronic techniques for tracing spammers: header analysis and honeypots. Our analysis suggests that spammers who understand identity-concealment techniques (or spammers who know how to use identity-concealment tools) are unlikely to be identified by a cybersleuth who relies on the electronic trail alone.

Note that we have not discussed future anti-spam technologies such as Yahoo Domain Keys or Microsoft's Caller-ID for email. The objective of such "authentication" systems is to render it impossible for a spammer to transmit email messages that appear to come from a reputable source. At this point it is not clear whether these technologies will eventually be deployed, nor is it known how they will a□ect spammer's behavior.

References

[1] Brightmail. Brightmail hails passing of U.S. Senate Bill 877. http://www.brightmail.com/pressreleases/102203_senate_bill_877.html.

[2] A. Curry. Proxypot. http://proxypot.org.

[3] Federal Trade Commission. National do not email registry, a report to Congress. http://www.ftc.gov/reports/dneregistry/registry.pdf, June 2004.

[4] LURHQ Threat Intelligence Group. Bobax trojan analysis. http://www.lurhq.com/bobax.html.

[5] L. McLaughlin. Bot software spreads, causes new worries. *IEEE Distributed Systems Online*, 5(6):1541–4922, June 2004. http://csdl.computer.org/comp/mags/ds/2004/06/o6001.pdf.

[6] A. Podrezov. F-Secure virus description: Agobot.FO. http://www.f-secure.com/v-descs/agobot_fo.shtml.

[7] V. Tanger. Multiple vendor HTTP CONNECT TCP tunnel vulnerability. bugtraq id 4131, Feb 2003. http://www.securityfocus.com/bid/4131/discussion/.

[8] P. Wood. Anyone for spam? *British Computer Society Review*, 2004. http://www.bcs.org/review04/articles/itsecurity/spam.htm.

[9] P. Wood. Save yourself from eternal spamnation. MessageLabs White Paper, April 2004.

Attachment A

Section 2.1 mentions that open proxies can be used for sending spam email that conceals the identity of the spammer. Here, I give more technical details on how these proxies are used. Spammers use two types of proxies: SOCKS proxies and HTTP proxies. We discuss each in turn.

Open SOCKS proxy. This type of proxy enables direct proxying of SMTP email traffic. As an overview, the spammer sends each email to a SOCKS proxy that, by design, forwards the email to its destination. This conceals the spammer's network address since the recipient only sees the proxy's address. The proxybench.com service currently lists 52 open SOCKS proxies that have been validated to function properly within the last month. In the discussion of bot-networks I mentioned that new SOCKS proxies can be constantly created by instructing bot-networks drones to run them.

Open HTTP and HTTPS proxy. Anonymous web proxies can be used for spamming in multiple ways. An anonymous web proxy is one that does not embed the network address of the client (e.g., web browser) in the request sent to the server. Because there are many such proxies available on the Internet, they are very attractive for spammers. The proxybench.com service currently lists 1,519 open anonymous HTTP proxies that have been validated to function properly within the last month.

The simplest way to exploit an open web proxy for spam is to use a web email service such as HotMail or Yahoo mail. These services embed an `X-Originating-IP` header in every email they send which completely identifies the sender's network address. Consequently, if the spammer were to send bulk email by connecting to HotMail directly the

29% of spam email messages that could potentially contain the spammer's network address within the headers, it is likely that a significant fraction was generated through bot-networks, open proxies, or untraceable Internet connections. The headers of spam email messages generated through most bot-networks, open proxies, and untraceable Internet connections will not contain any information about the spammer's network address. This means that the spammer's network address will actually be contained in only a small fraction of the 29% of spam email messages that could potentially contain the spammer's network address. And, as discussed previously, spammers use additional techniques (such as header spoofing) in such cases to mask which of the multiple network addresses contained in an email message's headers is the spammer's true network address. Thus, based on this informal study, I estimate that at least 71% and, more likely, as much as 90% of spam email contains no information identifying the spammer's network address, and is therefore untraceable. Below, I briefly explain the methodology used in my informal study.

This analysis is based on 10,468 spam emails received at my Stanford account in the 9 months from Oct. 2003 to July 2004. These emails were identified as spam by SpamAssassin. As discussed in Section 3.1, the study's goal is to bound the fraction of spam email from mail relays. Other identity concealing techniques (such as anonymous proxies and bot-networks) reveal nothing in the headers. The analysis focuses on the first Received header in each email. This header is inserted by my own email server. Let us separate spam email into two categories: (i) spam email for which the peer name provided in the HELO message does not match the peer's DNS name, and (ii) all other spam email. I found 7,472 email messages in the first category, which is roughly 71% of all spam email. Email in the first category is either a result of the peer lying about its name or a result of misconfiguration. For this informal study we discount the effect of misconfiguration. Now, spam email where the peer lies in the HELO message could not have come from a properly-configured mail relay. Indeed, properly-configured mail relays follow the SMTP protocol and need not lie about their DNS names. Thus, email in the first category is likely to be untraceable since headers will lead to a proxy or a drone, but will not lead to the spammer. Email in the second category could have come from a mail relay, but could have also come from a bot-network or an open proxy. Hence, 71% is a lower bound for the percentage of spam email that contains no information about the spammer in the headers. Because this is merely a lower bound, I estimate that as

much as 90% of spam is likely untraceable.

This informal estimate is consistent with an Oct. 2003 press release from Brightmail [1] stating, among other things, that "... 90% of spam is untraceable by available methods." Similarly, Wood [8] estimates that over 60% of spam email is sent through open proxies and hence is untraceable by header analysis.