Taughannock Networks

E-mail address forgery

A Taughannock Networks Mini-Paper
June 2004

In my roles as postmaster at CAUCE (the Coalition Against Unsolicited Commercial E-mail) and abuse.net, I get a lot of baffled and outraged mail from people who have discovered that someone is sending out spam, often pornographic spam, with their return address on the From: line. "How can they do that? How do I make them stop?" The short answers are "easily" and "it's nearly impossible."

One way that e-mail is very similar to paper mail is that you can scribble any return address you want on an envelope and mail it. With paper mail, just like e-mail, you can imagine ways to make it more difficult to scribble the name of someone you don't like, but the costs of doing so would be huge, and the benefits dubious.

For both paper mail and e-mail, it's not at all straightforward to determine who's allowed to send mail with what return address, nor from where people should be sending mail. With paper mail, I often drop mail from my wife in the mailbox, and occasionally from friends who've been visiting. Conversely, sometimes I mail my own mail, and sometimes the village clerk will send mail over my signature as the mayor. Sometimes I send mail at my local post office, sometimes I send mail from the other side of the country when I'm on a trip. All of these scenarios have e-mail analogies. Sometimes I send mail with my usual taugh.com address, but I also have addresses at AOL, Yahoo, Hotmail, Outblaze (another large free web mail provider that absorbed mail.com), `netscape.net`, professional societies such as `ieee.org`, and my college alumni association. I may have more addresses than most people, but it's quite common for people to have two or three.

When someone sends an e-mail message the return address is usually placed on the message by the user's mail program, such as Outlook Express or Eudora. The mail program then passes the message to a mail server, also called a mail transfer agent or MTA, usually provided by an ISP or company network manager. The MTA then sends the mail along to its destinations. For bulk mail, either legitimate or spam, the return address is placed on the message by a specialized bulk mail sending program. Some of those programs include the function of an MTA, while others pass the message to a conventional MTA for delivery.

One thing that's notably missing in this process is any kind of security. The user's mail program or bulk mail sending program can use any return address it wants. This may sound like a bad idea, but the reality is that only the user (or the person running a bulk mailing program) knows what addresses he's allowed to use.

Taughannock Networks provides strategic consulting, design, and prototype implementation on e-mail and Internet infrastructure. +1 607 330 5711 or info@taugh.com

Some ISPs have attempted to verify the addresses on mail going through their MTAs, with little success. Bell Atlantic, a predecessor of Verizon, used to require that all outgoing mail through their MTAs had an address at `bellatlantic.net` or one of the other domains of ISPs they'd absorbed. This technique turned out to be both annoying to their users and useless to prevent spam. It was annoying because all of the users who had valid addresses elsewhere couldn't send mail with those addresses, and it was useless because their system wasn't able to tie a particular address to a particular PC, so spammers merely made up fake `bellatlantic.net` address and spammed away.

For Internet e-mail's first fifteen years, address forgery wasn't a problem. Technically it was easy, but there was little incentive to do so, and it was rare other than as a prank. In recent years, spammers have put forged addresses on most of their spam, both to try to defeat filters, to make it harder for recipients to figure out where to complain, and occasionally to annoy the legitimate owner of the addresses. For a while, spammers made up random addresses, but as recipients started filtering out mail with non-existent domains in the return address, spammers adapted by using real addresses, often taken from the same lists as the spam targets. A related but separate problem is *phishing*, impersonating a trusted organization to trick people into revealing financial information.

In the past year there's been a great deal of work trying to figure out some way to deter address forgery. It would be straightforward to invent a system that registers a single mail source for every Internet domain, and require that all mail from a domain come from the registered source. While that would be very useful for some domains like `paypal.com` that are often forged and already send all their mail from one place, it would break a surprisingly large amount of legitimate e-mail, from e-mail discussion lists to electronic greeting cards to automatic mail forwarders. Several validation schemes are in the works, with names like SPF, Caller ID for E-mail (those two recently merged) and TEOS, and Domain Keys. But it remains to be seen both whether such schemes can work with the many legitimate but unusual mail sending methods that they don't easily cover, and more importantly whether spammers will just find ways to send their spam with valid domains. The majority of spam is sent through virus controlled "zombie" computers, so the spam could easily forge the zombie's own domain. Or since spammers already register large numbers of domains, they can use those domains in their spam and publish validation rules that the spam satisfies.

The whole issues of on-line identity, forgery, and authentation are remarkably complex, so we don't expect any resolution to the forgery problem soon.