

## Spamming, Phishing, Authentication, and Privacy

Steven M. Bellovin

It isn't news to anyone that email is becoming almost unusable. Unsolicited commercial email (spam) peddles a variety of dubious products, ranging from pharmaceuticals to abandoned bank accounts. The so-called "phishers" try to steal user names and passwords for online banking. And then, we have viruses, worms, and other malware. Although there are would-be solutions to these problems, some of them have the potential to do far more harm than good.

An obvious approach to spam-fighting is some form of sender authentication. If we know who *really* sent the email, we can deal with it: either accept it, because it's from someone we know, or—if it's spam—we can chase down whomever sent it and take some sort of judicially-sanctioned revenge. It's simple, it's obvious—and it doesn't work, for a number of reasons. Fundamentally, most people accept—and want to accept—email from more or less anyone. Just while writing this essay, I received no fewer than five legitimate pieces of email, addressed directly to me, from new correspondents. It does no good to have some assurance of a total stranger's identity if you're going to accept the email anyway. Fundamentally, identity is only a concept that makes sense within a shared context — without which the sender's authenticated identity, as opposed to the merely asserted identity, means very little.

Of course, spammers can authenticate themselves, too. Just as today they buy throw-away domains, in a world of authenticated email they'll buy throw-away authenticated identities. Indeed, anecdotal evidence suggests that the spammers have been the fastest adopters of the prototype authenticated email schemes. We thus have the following conundrum: if you use these anti-spam techniques, statistically you're *more* likely to be a spammer!

Beyond that, remember that much spam comes from hacked machines. Someone who "Ownz" your machine can steal your online identity quite easily, including (of course) any cryptographic keys you possess.

If authentication techniques don't work against spam, do they help protect us from phishers? Here, at least, there is reason for optimism: a phishing attack *is* an impersonation attempt; if we can really authenticate the sender of the email, would we not be safe?

Unfortunately, the proposed email authentication techniques won't do the job. What you really want is proof that "this is the party to whom I gave my money"; all

that this scheme can establish is that the sender owns some plausible domain name. It says nothing about your prior relationships. We don't have to imagine this attack; one of the very first phishing incidents involved email appearing to be from paypal.com; the actual domain was paypal.com.

Consider, instead, a scheme where, when you opened an account, the bank sent you a copy of its certificate. This certificate could indeed be used to authenticate any email from the bank. Note the crucial difference: such a certificate is bound to a previous transaction, rather than to a name.

Authenticated email solves some problems. If nothing else, it hinders the spread of email worms, since the infected machine will be positively identified. Furthermore, there are some situations where a list of permitted senders is in fact used. In the best of these schemes, purported identity is used to drive some sort of challenge/response scheme. Authenticated email would provide some protection here, though even without it successful "joe jobs"—forgery of a legitimate user's identity—are relatively uncommon. The spammer would have to select source-destination pairs of addresses to bypass simple-minded permitted sender lists.

However, there are serious disadvantages. Some are logistical: with some of the proposals, inbound mail-forwarding services such as acm.org won't work properly; people will be sending mail from more or less anywhere that claims to be from acm.org. Other schemes have trouble with mailing lists, such as those that add administrative information to outbound messages.

But the most serious problem is one of privacy. If all mail *must* as a practical matter, be signed, all mail becomes traceable. (Many anti-spam payment schemes share this problem.) The U.S. Supreme Court has noted that "anonymous pamphlets, leaflets, brochures and even books have played an important role in the progress of mankind. Persecuted groups and sects from time to time throughout history have been able to criticize oppressive practices and laws either anonymously or not at all . . . It is plain that anonymity has sometimes been assumed for the most constructive purposes." Do we want an electronic world without such advantages?

**Steven M. Bellovin** is a researcher at AT&T Labs in Florham Park, NJ.