Stop, in the Name of Spam

I READ THE ARTICLE BY LORRIE Faith Cranor and Brian A. LaMacchia ("Spam!," Aug. 1998, p. 74) with interest. I find the title is inaccurate; I still use the old definitions of the terms "hacker" and "spam," and I insist that people use the correct terms of "cracker" and "junk mail" every time I see or hear them misuse these terms.

Nevertheless, as the former senior Internet mail systems administrator at America Online, and the primary person to implement anti-junk mail techniques on behalf of AOL users, I still have a rather unique perspective on the problem that few other people in the world can appreciate—they just haven't sat there and personally watched millions of junk mail messages flood their systems on a daily basis. Furthermore, in a matter of *hours* I've seen junk mailers work around new protections that took me weeks to devise and safely implement. So this is an arms race few people outside AOL have seen, or are perhaps capable of appreciating.

With regard to the survey data, I have more recent estimates (within the last couple of months) from knowledgeable sources that indicate they are seeing complaints from their customers on the order of 30%–40% of all mail being junk mail, and on bad days, these sites are seeing 50%–60% junk mail. Obviously, the situation has gotten much worse since the cited surveys.

One of the biggest problems with junk mail is people using throwaway dial-up accounts to generate the millions of messages. I am currently working on co-authoring an RFC that will address some of these issues.

One solution involves getting service providers, such as UUNet, MCI, AT&T, and AOL, to put their dial-up lines in a separate subnet from the rest of their machines and to publicize that information. Furthermore, they should ensure proper reverse DNS is set up correctly for all the IP addresses assigned to dial-up customers and put those machines in a different subdomain, separate from the other machines.

These two steps alone would allow sites around the world to refuse to accept port-25 connections from a dial-up customer (under the theory that if you're a legitimate customer of the service provider, you should be using the Internet mail services it provides for all legitimate customers and not attempt to connect to an external site directly). Some of these service providers are already doing this, at least to a limited degree; AOL has .ipt.aol.com and UUNet has some dial-up lines in .da.uu.net.

I would extend this by having the ISPs themselves cut off external access for their dial-up customers to the outside world. If these dial-up customers wish to have Internet mail delivered, that mail should be transmitted to the Internet mail servers the service provider offers, and once it's out of the customer's hands, it should be out of their minds as well.

There may be some legitimate customers who believe they have a right to directly access the Internet mail servers of the outside world, and the service providers should take this as a sales opportunity to offer them a higher-cost-level of service providing such access. But, that also requires a higher level of information investment from the customer, and a higher level of penalty should the service provider discover the customer is

PAUL WATSON

using this access to transmit or relay junk mail.

If we implemented these suggestions, starting with the large service providers, we should quickly close most of the easily obtained throwaway accounts used for generating junk mail and reduce junk mail traffic from the untold multitudes of individual users sitting at home with a single PC and dozens of dial-up accounts.

Instead, we'd be left with companies like Cyber Promotions, and we can more easily deal with them by cutting off their Net access or simply refusing to accept any mail whatsoever from any site known to transmit their mail.

I further suggest that if the backbone service providers of the world band together and decree that junk mail is anathema to their existence, and any site that generated or relayed junk mail would have to stop or have their network access cut off, that would solve the overall problem.

The Cyber Promotions of the world would be completely shut down, because no ISP would be stupid enough to give them access, lest they have their own access severed.

This is the FidoNet solution to the problem—don't be excessively annoying or excessively annoyed, lest we cut you off at your upstream feed, and if your upstream feed doesn't cooperate, we'll cut it off from its upstream feed, ad infinitum.

Note that this also solves the problem of individual users generating junk mail, since no ISP in its right mind would take any but the strongest possible measures against customers that were to violate this rule, lest they lose their very livelihood.

Brad Knowles
Bethesda, MD

On a personal note, spam is a painful phenomenon. I was thinking that spam's anonymous aspects, due to forged headers, is what makes it difficult to enforce accountability. This is confirmed by the notion of attacking spam's anonymous nature through law.

What if all mail servers were registered and required to generate private and public keys in an antisymetric cryptographic system and to prepend the header, as well as a verification field. This field would contain a message digest for the header and another for the body of the email. The verification field would be encrypted with the private key.

If the public keys were published, I could at least be sure the message came from where it claimed to be sent from. I could easily filter out any message not having the verification or having an invalid one where one or both of the message digests didn't match.

What about those ISPs that welcome spammers? If they failed to take corrective action against spamming complaints, those ISPs would be banned.

Free speech arguments for spam are hogwash. Not only is it commercial speech, but in any other media I have the choice *not* to listen. I can walk quickly by the radical on the soapbox in the park. I can turn off the radio and TV, or change channels. I can tell the telemarketer on the phone, "I'm not interested and please remove me from your list." Email forces me to wade through a sometimes lengthy list and sometimes waste a lot of download time. Spammers' free speech rights cannot override my right not to listen or my economic right not to subsidize their speech.

Dan Bates
Lewisville, TX

I eagerly looked forward to reading the spam article when the latest issue of *Communications* arrived. However, it was a letdown. The writers failed to take a firm stand on how to address the problem.

Consider past mechanisms of mass media: the U.S. Postal Service regulated junk mail and its contents; newspapers and magazines have always regulated advertising content; and radio, followed by television, did the same as these mediums matured.

Telemarketers, then, took to the phones and incorporated high-pressure sales techniques. Like other media, these soon required regulation to control what was happening, including the ability, under penalty of fines, to be removed from calling lists.

When fax machines came along, it again took federal regulation to protect this method of commerce from being destroyed by marketers.

Now, with email, comes spam. Spam is destroying the utility of email and will require regulation. Corrective regulation will control identification of senders, content of the advertisement and sales material, and provide penalties for failure to comply. Otherwise, spam will continue unabated. If we direct the cost of sending spam back to the spammer, then spam will be legitimate and affordable to those with the

money to force junk mail into our email containers.

The writers mention that sending spam is inexpensive. Wrong. It is expensive to the carriers and receivers. That expense must be pushed back to the senders.

I believe spam must be banned. Why is it that every method of reaching the masses must be burdened by marketers? Let's keep email free and clear of marketing and leave it as a pure communication vehicle. It is not necessary to turn every effective communication mechanism into a marketing battleground. Opt-in should be the legal, required default. If marketers do not have a signed legal release, then they can't send you junk email. Hard line, yes. The best thing to do … absolutely.

Ken Kashmarek
Eldridge, IA

I'm personally opposed to any additional legal or government regulatory approaches to spamming. Spammers who commit fraud or break other existing laws should be dealt with according to existing remedies.

The most important technical problem to overcome is that of identification. We must be able to trace mail back to its source, thus raising questions regarding anonymity—issues and concerns that are analogues with telephone caller ID.

I propose that all anonymous mail be marked in an unambiguous and machine-readable fashion. ISPs could then offer services allowing anonymous mail to specific addresses.

Another option would involve an enhanced mail protocol with an optional three-way handshake. My robot gets your initial mail, hashes your reply address with a secret, and sends a reply. You have to resend your message with the correct cookie to get past my robot. This way, I've assured that mail I receive is from some sort of valid address.

Another avenue of accountability to supplement this model (since one could easily create an account, use it to spam, and then abandon) is the mechanism that prevented spam for the first 15 years of the Internet's existence.

In the old days, when the Internet had, at most, a few thousand hosts, peer pressure prevented spam. If I, as a system administrator, was lax in controlling the behavior of my users, I'd be cut off. Since the Internet is a network of autonomous cooperating networks, any site can block any other site—and a bad site will be blocked by a large number of sites. The problem with informal methods is that they don't scale well to the modern Internet. System and network administrators in this environment usually don't have the authority to set such policies.

However, we can apply formal methods and new protocols to cooperatively scale this peer-conformance model up to contemporary requirements. Some of the model's elements are already being implemented in "sendmail" and other mail transport agents (MTAs). For example, it is now standard practice for a receiving mail service to do a "double reverse lookup" to validate correspondence between a sender's (or relayer's) IP address and its claimed domain name. Discrepancies are noted in the headers of the received mail items, and sites can (optionally) refuse to accept mail from unauthenticated senders.

Conventional practice used to allow promiscuous relaying. My mail host would accept mail from anywhere and relay it as a courtesy. This was common practice primarily because it was simple to implement and easy to maintain, and there was little or no perceived risk or cost associated with it. However, spammers have abused it—forcing most modern mail administrators to add and maintain additional configuration information to their MTAs. This ensures that each mail hub receive only mail destined for one of its approved client domains.

Another approach is Paul Vixie's Real-time Blackhole List (RBL)—an elegant combination of reverse DNS and MTA relay authentication that allows one to maintain a blacklist of spammers, sympathizers, and collaborators.

The beauty of this approach is that it benefits from the scalability and performance of the existing DNS infrastructure. The responses are even cached in the same way as any other DNS entries.

The downside is that many addresses are dynamically allocated. Thus, we hold an entire ISP or organization responsible for abuse by any of its subdomains. But it is reasonable that the ISP should also be responsible to the rest of the Internet for failure to enforce reasonable acceptable-use policies.

Internet participation entails the acceptance of a social contract—one that cannot tolerate spam and other forms of abuse. We would certainly disconnect a site that deliberately propagates invalid routes (in effect, stealing the traffic that was destined to other peers). We should do the

same for other forms of abuse.

The main problems with Vixie's existing scheme are that new records must be added and removed manually, and U.S. liability and antitrust laws and conventions may prevent this method from being sponsored or maintained by a private entity.

I propose establishing a number of cooperative groups. Each would have voting members (the system administrators subscribing to it). Each member would have a digital key and would vote using that key. Any time a site was held by a member to be abusive, that member would vote on the addition of that site to their RBL for a given protocol (mail, news, and so forth). Past a set threshold, the rogue site would be added. Getting off of that list would require a number of revocations and opposing votes.

Thus, a rogue site would rapidly get its IP addresses locked out of large numbers of sites. The process of getting those IP addresses removed from the blacklists would be time-consuming enough that ISPs would not allow their customers to cause such a problem.

The technologies are certainly out there. Engineers and mathematicians far more qualified than I could work out the details of a future Internet model.

Otherwise, spam *will* make email and Net news unusable, and will rip the fabric of the greatest cooperative communications venture in human history. We cannot allow that to happen.

Jim Dennis
Campell, CA

Please address all Forum correspondence to the Editor, *Communications*, 1515 Broadway, New York, NY 10036; email: crawfordd@ acm.org.