;login:

THE MAGAZINE OF USENIX & SAGE

April 2004 • volume 29 • number 2

inside:

USENIX

The Advanced Computing Systems Association

musings

by Rik Farrow

Rik Farrow provides UNIX and Internet security consulting and training. He is the author of *UNIX System Security* and *System Administrator's Guide to System V*.

rik@spirit.com

In one of those weird twists of fate, being popular meant that you got bombarded with many copies of the MyDoom virus in late January and early February 2004. The MyDoom virus (A version) used its own SMTP engine to mail copies of itself, as an attachment, to names culled from users' address books. The same email addresses were also used as sender addresses, so you would also get besieged with bounce messages. If you were deluged with copies of MyDoom, you just received a measure of your online popularity.

Adding insult to injury, most anti-virus email server products, when armed with a signature to detect MyDoom, would send a notification to the forged sender, so you were also bombarded with assertions that you had sent the MyDoom virus, and thus, were infected. By implication, you should have bought the AV vendor's product so you could be protected, and spam other people by sending messages to forged sender addresses.

You might think that I am being unfair, but a simple inspection of mail headers would show that the alleged sender of MyDoom had nothing to do with the source address of the system that sent the email. You would think that companies that deal with viruses, which get spread mainly by email, would at least be capable of running trivial checks on email headers before spamming people with useless notices.

MyDoom turns out to have another interesting tie into spam. Some security friends were discussing the virtues of blacklisting entire netblocks. Many people have already decided to reject all SMTP connections coming from Asian netblocks, as many of these countries have large broadband (mostly DSL) networks of home users, and these home systems have become a popular haven for the open relays used by spammers. Some of my friends were going further, deciding that they should also be blocking the netblocks used by the large US-based broadband providers, such as Comcast, ATTBI, and RoadRunner.

To me, this appeared to be a little extreme. So I decided that I should take a deeper look. I started by grepping the Received lines added to incoming email by my own email server. My Spam directory had, at that time, about 4000 emails in it. I also asked Brian Martin, a friend from attrition.org who not only saves spam emails but diligently sends messages off to abuse@whatever for his Received headers. Together, we had close to 10,000 Received lines that included both the sender's IP address and the domain name resolved by Sendmail or postfix in the form

```
Received: from mail11.cybertrails.com
    (mail.cybertrails.com [162.42.150.35])
```

The system that delivered this particular email to my system gave its name in the HELO message as mail11.cybertrails.com. The source IP address during the TCP transaction was 162.42.150.35, and the reverse lookup of that address was mail.cybertrails.com. Note that I am not beating up on Cybertrails, a past ISP that apparently continues to forward spam to me based on an email address that I never used and have asked them to kill. They are now on my REJECT list.

Next, I used a Perl one-liner to extract the resolved domain name and IP address from the Received lines, then used another one-liner to extract just the last two names in each domain (cybertrails.com in the above example). Finally, I ran that result through

a pipeline that sorted the names, used uniq -c to count them, and sorted the counts by uniq in reverse order. Now I had a sorted list of the sources of the 10,000 spam messages that Brian and I had received over about a month.

The results were revealing, but not too surprising. Comcast.net and attbi.com together accounted for 15% of all spam received, with RoadRunner coming in third with 5%. Many other well-known broadband providers fit into the top 20 offenders, for a total of over half of all the spam we had received.

Now, this is not a scientific survey, just a quick peek at a sample of spam. I make several assumptions, such as that I trust Brian's and my own SMTP MTA to correctly resolve addresses, and that spammers are not adding forged Received lines, complete with domain names and addresses in the format used by UNIX MTAs. But it appears to me that my friends who have been considering blocking all email from broadband providers as a way of stopping spam really do have a good point.

One person in the group, known as Hobbit, has already blocked many broadband providers, but with a very reasonable twist. He permits SMTP from the mail servers for each broadband domain, based on the MX records for that domain. I thought this was a great idea, as it means that only email relayed by the SMTP servers for a domain would be permitted. Most spammers will not use an ISP's mail relays, but send email directly, using their own SMTP engines installed via viruses or other attacks, or via open relays. Hobbit could still receive legitimate mail, relayed from "registered" servers – those that have MX records for the sender's domain.

When I started looking deeper, I discovered that what Hobbit had done was a lot of work. For example, there is a *fl.comcast.net*, but fortunately only two MX records for all of comcast.net, corresponding to four IP addresses. Other broadband providers actually support a set of SMTP servers, represented by MX records, for each subdomain, which are often organized by state. And this information is subject to change.

I decided I didn't want to maintain a comprehensive list of networks to block while permitting only registered SMTP servers within these blocks. Sure, blocking some netblocks looks like a good idea, something that would quickly free up my system for other tasks (using client_acl of Postfix or your favorite MTA to reject connections from clients that don't correspond to MX records). But I think that there is a better way.

What if ISPs blocked outgoing connections to port 25 that come from clients that don't have their own domains and MX records? Some ISPs do this already, although obviously not some of the larger ones. Doing so would require that the ISP actually maintain filter rules that would permit the SMTP servers that it knew about, while blocking all others. It appears to me that doing this would block the most commonly abused spam relays. It would also prevent the ISPs from being added to spammer blacklists.

Of course, getting ISPs to do any useful filtering has been wishful thinking so far. Back in 1998, RFC 2267 sought to prevent source address spoofing by encouraging ISPs to block packets that enter their networks from leaf networks that have spoofed source addresses. Source address spoofing on the Internet would have vanished if this had been accomplished. Obviously, it hasn't.

Now what do you think would have happened to MyDoom, and SoBig, and many other of the big viruses that spread using their built-in SMTP engines, if these filters were in place? The ISPs' filters would have blocked attempts to connect to any SMTP servers outside of their networks, slowing the spread of these viruses to a crawl.

ISP filtering certainly could be done. One of the reasons it is not done is that it means more work for the ISPs – work that many are loathe to do, as it costs money and affects the bottom line.

So we get deluged with viruses, virus bounces, AV warnings, and spam, largely because of ISPs not willing to take responsibility for managing their own networks properly. Perhaps it is time for a little regulation.