

The FTC Has Been Duped by Radical Antispammers

Executive Summary

The FTC should reconsider its view on open relay operation as it relies on several invalid assumptions. Several valid technical reasons for open relay operation do exist. The Commission should open the dialogue on open relay operation to include open relay operators; the composition of the advisory group is currently unbalanced in favor of radical anti-spammers. The FTC is also blocking email from open relay operators, and thereby excluding them from participation in public process. I urge the FTC to withdraw the "Facts for Business" release entitled "Open Relays - Close the Door on Spam", and to investigate the false and misleading claims made by certain open relay blacklists to the FTC and to their customers.

Detailed Analysis

My name is Dean Anderson. I own Av8 Internet. I have run open relays since 1996. It has therefore been my business to know who is abusing our open relays.

The italic'ed quotes are taken from the FTC "Facts for Business" release entitled "Open Relays - Close the Door on Spam"

"Not only could this overload your server but worse still, it could damage your reputation because it will appear that you sent the spam".

Antispammers frequently make these claims, but they aren't true. While it could appear so to an extremely naïve user, it would not appear so to any sophisticated user. All relays (open and closed) insert "Received:" headers in the message. These headers indicate the source IP address of the abuser. The header inserted by the Relay cannot be altered or removed by the (ab)user. Thus, while the abuser can forge other headers, such as the

from: address, and can add additional forged "Received:" headers, they cannot hide their IP address. Thus, they have not hidden their identity, or truly made it appear that the mail was legitimate.

As for the other forged headers such as the from: address, these headers can be forged whether or not they use an open relay. The open relay did not contribute to, or enable that behavior. They can do that anyway.

"There is little, if any, benefit to you in allowing this email to cut through your server; the people in your organization aren't receiving or sending it"

Another false claim, frequently heard from radical antispammers. Open Relay is the only method to send RFC 821 SMTP protocol email from IP addresses outside your IP address space. When my mail customers get connectivity from another ISP, they are accessing our mail server from that other ISP's IP address space. Therefore, this is an open relay transaction. Assuming they send mail to someone not on our service, both sender and receiver are outside our IP address space. These users are our customers. Clearly, there is a benefit from allowing our customers to use our relays from other ISP's address space.

There is an alternative protocol proposed to handle this situation, called SMTP AUTH. It was invented to compensate for the failures of a previous scheme, called POP-BEFORE-SMTP, which required impractical, non-standard, back channel communication between the POP server and the SMTP server. The SMTP AUTH protocol has not yet been standardized, and is not widely supported by email clients. There are something like 12 or 15 clients that support it. Notably, popular clients such as Microsoft Outlook, and Qualcomm Eudora, and Netscape Communicator support it to varying degrees. However, about a thousand or so other clients do not. Most have no plans to offer support in the future. This protocol has a number of flaws outside of its lack of standardization. It works fine for simplistic mail situations, such as found in a residential dialup ISP, or sometimes a small company. It does not work very well if the email situation is more complicated, such as when an insurance company has 50 sites spread around the country,

and is using a dozen or so ISPs to provide connectivity to those offices. Supposing it has its own corporate mail server, and want to purchase backup SMTP service from an ISP, that ISP would have to either create and manage thousands of accounts for their employees, or run an open relay server. Clearly, the open relay server is the right choice.

Another flaw is that the user has to maintain possibly separate usernames and passwords for sending and receiving email. A user that has outsourced their mailbox to one ISP, may still need relay service from another ISP. Thus both ISPs have to maintain usernames and passwords for those users, who may not get any other services from the ISP. In the case of a residential dialup, this isn't a problem. The user will already have a username and password. But in the case of a large company that has thousands of users, the ISP generally doesn't have usernames and passwords for each employee. The company manages that information. There are other flaws to SMTP AUTH that I won't go into here, as I think I have made the point. However, these are not hypothetical situations. Av8 Internet (and other ISPs) have a good number of customers with such outsourced services.

After all this work, SMTP AUTH has not reduced spam. A spammer is always the customer of some ISP. They will therefore have the SMTP AUTH password, and can still send mail. If the abuser has this username and password, the abuser can do anything they can do with an open relay, or a closed relay. They have all the same capabilities they had before SMTP AUTH. Nothing has been changed with respect to abuse issues. The ISP will not learn of this until someone complains about the user. The complaint will of necessity indicate the date, time, and IP address of the abuse. This is enough for the ISP to determine which user conducted the abuse. SMTP AUTH has not contributed anything to the cause of reducing spam.

Theoretically, it is impossible to control spam with a protocol. This follows from work done by Claude Shannon, in which he proved that one couldn't prove the non-existence of a covert channel. Applying this to email, we cannot prove that abuse (a covert channel) cannot be sent (the channel doesn't exist). Thus, no email protocol can be free from

abuse. This is completely obvious once you realize that all spammers are the customer of some ISP somewhere. Radical antispammers always couch their arguments so that it seems that spammers are somehow outsiders that have been let in (by open relays). This isn't true, which is why efforts like SMTP AUTH, and its predecessor POP-BEFORE-SMTP are doomed to fail.

"In the early days of the Internet, many mail servers were kept open to allow email to travel among different networks."

This is myth. But it is enlightening to understand the actual history, and how this myth evolved. In the early days of the internet (1986), it was thought that reverse DNS could be trusted to identify hosts, and that it was "silly" to put the numeric IP address in the Received: header. This message from Rob Austein to the TCP-IP list illustrates:

Message from Rob Austein (SRA@XX.LCS.MIT.EDU)

Wed, 1 Oct 1986 16:00 EDT

=====

Date: Tuesday, 30 September 1986 13:29-EDT

From: The lost Bostonian <gds@spam.istc.sri.com>

To: header-people@mc.lcs.mit.edu, tcp-ip@sri-nic.arpa

If it is true that all IP implementations enable a server program to determine the IP address of its peer, then the HELO command, and its response could be eliminated, which would save us a few bytes.

You are assuming that it is always possible to translate addresses to names and vice versa. Unfortunately, there are some people out in the

world running domain nameservers who are totally clueless about what they are doing, and there are others who have the misfortune to be stuck behind a losing gateway or otherwise be unreachable much of the time. Do you really want to make it impossible to receive mail from some host because a third party is broken? **Or have to put numeric addresses into the Recieved headers?**

The answer is to fix the silly net, not throw away features to save two IP packets.

--Rob

Unfortunately, Rob's view prevailed: No one put numeric IP addresses in Received: headers. The result was that if the client machine didn't have reverse DNS, or could fake the reverse DNS response, then it could send totally anonymous and untraceable email. This is significant, and is the source of several myths about open relays.

This condition persisted until shortly after the Internet went commercial, and serious, untraceable abuse began to appear. This is now called "anonymous relay" to distinguish it from "open relay". The solution to this problem was two-fold: The Received header was changed to include the numeric IP address, and it was urged that unnecessary relays were closed if they weren't needed. Of course, until the software was updated, closed relays were still vulnerable to this sort of abuse.

The problem was made more severe by the fact that early on, the Internet was dominated by Unix workstations, and that every workstation had a mail server running, regardless of whether it received email via SMTP. So it was trivial to exploit this problem. It was much easier in most cases to disable the mail server (close the relay), than to upgrade the software. By 1996, there were few if any such buggy SMTP servers still running, and the

anonymous untraceable email exploit was no longer possible. What started out around 1993 as a reasonable request to close unnecessary relays had evolved into a baseless and unreasonable crusade by 1997, when I first encountered the crusaders.

Email abuse was not the only harm done by inappropriate faith in reverse DNS. The Morris Worm which shutdown the internet in 1988 also exploited a weakness in the BSD r-command suite, due to reverse DNS.

It is a historical irony that many of the same anti-spammers who parrot the harms of open relays, also continue to promote the inappropriate use of reverse DNS, apparently still trying, as Rob lamented: "to fix the silly net". Even though the Shannon's theorems stand against their efforts, they persist. Tilting at windmills seems to be their favorite pastime.

It is also a historical footnote, that the HELO command is frequently used to create a false impression. It probably would have been better to have listened to the "lost Bostonian" in 1986 (quoted above), and eliminated HELO from SMTP. It is fair to say that Rob Austein (unintentionally of course) created the opportunity for Spam in 1986.

It was never the case that in the early days, open relays were kept open "*to allow mail to travel between different networks*". Assuming this is a reference to the operation of relays which translated between different network protocols, such as Usenet UUCP, and BITNET, and so on, it is still wrong. Those relays were close-able (if not closed), since mail destined to UUCP would still be locally delivered, so far as SMTP was concerned. By definition, an Open Relay allows mail to flow between two different IP addresses, which don't belong to the Server operator.

"But today, an open relay is most likely to be used by a spammer. Using automated software, spammers scan the Internet for an open relay"

This isn't so much a provably false statement, as it is just a statement that varies from my extensive experience operating open relays. Some History is in order:

Unlike many ISPs, I have been blessed with a larger than usual block of IP address space. I have a /16 (65000 IP address) and 2 /21's (4 thousand each) for a total of around 73,000 IP addresses. The /16 is where much of the unused space was. I have been associated with this /16 almost since it was assigned by the SRI back in 1988. In 1989, I worked for the company that it was assigned to. I left that company in 1991, and came back as a consultant from 1995 through 1998. Through contract with that company, I took over administration of these blocks. Some parts of that address space has never, ever, been used. I've been in a position to know which parts have been used and which parts haven't.

I began operating open relays in 1996. In 1997, I responded to a post on a mailing list made by someone claiming that all open relays should be closed, and asserting that there were no legitimate purposes to open relays. I responded, explaining what the legitimate purposes were. We were blacklisted by group called ORBS.ORG. A few days later, we got our first open relay abuse. And I made my first relay abuse complaint.

Later, I began to log, using a Cisco Access Control List, the IP address of everyone trying to connect to port 25 (SMTP) in our address space. As we had a lot of unused addresses, connections to an unused address indicated scanning. It was then a simple matter to check the mail server logs to see if that IP address tried to send mail. Sometimes they did. In order for a scanner to learn if a server is an open relay, they have to send an email through it. So by looking in the server logs to find the email addresses from the message, I could find the email address of the scanner. That email address had to be real, and the scanner had to have access to that email addresses' mailbox. What I found was that the scanners were open relay black lists. After a time, I discovered that only open relay black lists were scanning for open relays.

To date, only one company besides an open relay blacklist has ever scanned our address space: Rockliffe.com. They make an email server. I sent a complaint about their unauthorized scanning, but they didn't respond or explain their purpose. Their scan took 3 days to cover the entire list of 65000 IP addresses.

So next, I setup a open relay mailserver on a previously unused IP address. I had the logs, so I knew that this IP address hadn't been recently scanned. This mail server wasn't used by anyone. I then submitted this address to an open relay blacklist. It was promptly scanned by the blacklist. A few days later, the server (without being scanned by anyone else), began getting relay abuse by spammers. Interesting! Spammers had learned of the open relay through the blacklist! So I closed the relay, and had it re-tested. Several days later, the connection attempts (from my logs) began to drop off. Spammers had also began to forget about the relay. Testing again, with a different blacklist and a different IP, I observed similar behavior, but different client IP addresses. So I posted an indictment of the open relay blacklists to a mailing list. I said that the blacklists were promoting spam by advertising IP addresses of open relays. I said I could tell which blacklist the spammer got the IP address from.

Immediately after that, the spammer behavior changed. They began to "pre-scan" the /24 (255 IP addresses in a block) surrounding the relay. They only connected to addresses, but didn't send mail. My test IP addresses had been in the same /24. Now I couldn't tell any more which blacklist was promoting the abuse. However, I was definitely in communication with the spammers, as I my post had caused them to alter their behavior. I suppose it isn't strange that spammers would monitor an anti-spam list. However, this list had only about 1000 subscribers. So there couldn't be very many spammers. And why would spammers care if I knew which blacklist they were abusing???

At this point, I considered the open relay blacklists to be something like "useful idiots" who were doing things that spammers could exploit, by giving out lists of open relays. I did not think they really intended to help spammers.

Next I tried to get the blacklists to list IP addresses that couldn't be abused by spammers. This would make it impossible for spammers to abuse the blacklists, to find relays. I knew that not many people used the blacklists, and I didn't much care and couldn't stop them from listing our addresses. But I could stop those addresses from being abusable by

spammers. So I separated the relay's input address from the output address. The input address is the one given to the customer. The output address is the one that it uses when it connects to other mailservers. However, the output address won't accept incoming connections to port 25. So if the spammer gets the IP address from the blacklist, it won't be able to send mail. To my chagrin, some blacklists started trying to detect (and then advertise) the input addresses. There is no bonafide use to this input address list. It was useless except to spammers. Any mail sent from our relays would come from the output address. Anyone, who wanted to block open relays, would have to block the output address. Blocking the input address would have no effect. The blacklists were actually going out of their way to help spammers abuse open relays. Interesting! They weren't simply "useful idiots," they were active participants!

In another incident, we had a anti-spammer overtly abuse our relays. His name was Christopher Neill, and he was an abuse admin for Verio. During one of these mailing list exchanges, he began to abuse our relays. Very overtly: "Hey Dean, see what I can do -- Chris". That sort of thing. I sent a complaint to Verio. This continued. I sent more complaints. Eventually Verio fired Mr. Neill, as he reported to the SPAM-L list. In his own words, he blamed me for being fired. According to Chris, the FBI investigated him. Evidently, Verio's management and legal staff must have disagreed with his assertions that open relays were free. There have been other incidents where people have set out to "teach me a lesson" by abusing our relays.

Since getting good at blocking the open relay blacklists from scanning, our relay abuse has dropped to nearly zero. So, in my experience, it is more likely that anti-spammers will abuse open relays.

This may not be the experience of every open relay operator. Most open relay operators won't admit they have open relays, and certainly won't confront any anti-spammers about the legitimate purposes of open relays. But I can't take back the statements I made in 1997. Perhaps my acts of being a spokesman for open relay use have made me a target for anti-spammers.

Clearly every organization has a few miscreants. For example, I was President of the League for Programming Freedom (LPF), which was founded by Richard Stallman, who also founded GNU. When the LPF boycotted Lotus and Apple Computer for the assertion of User Interface Copyrights, I had some people come up to me and offer to take a job with Apple and sabotage their products. Our lawyers (professors of Law) had by then advised me about anti-trust. I knew that not only was this illegal, but that the LPF could be responsible for any damages. Anti-trust makes it illegal for group boycotts to harm legitimate business. (Anti-trust applies to blacklists, too. See *Exactis V MAPS*). I had to tell those people in no uncertain terms that if they did that that we would immediately inform Apple and eject them from the LPF. No doubt, there are anti-spammers who might try similar activities. Perhaps my speaking out has made me a target of those people. But one would think that some real spammers would try to abuse our relays. To date, we have not observed this. What is unknown though is how many such anti-spammer miscreants there are, and what effect they have on things like total open relay abuse.

"If they find your server is open, they route their bulk mail through it, spamming in greater volume and less time than they could using their own individual computers"

This is frequently asserted by anti-spammers. However, it is not true. Having operated open relays for 7 years, I know what happens, and what abuse runs actually look like. When a message is sent via SMTP, the user mail agent sends recipient list, and then the contents of the message. As follows:

```
HELO some_string_just_made_up
MAIL FROM: <someone@someforgeddomain>
RCPT TO: <recipient1@domain>
RCPT TO: <recipient2@domain>
RCPT TO: <recipient3@domain >
...
```

RCPT TO: <recipientMAX@domain >

DATA

Forged message headers

Text of message

.

The lines with MAIL FROM and RCPT TO are called the envelope addresses. These are the addresses the message will be delivered to. If the @host part is the same, the open relay will just connect to the mail server for domain, and send exactly the same sequence of commands. Except that it will insert its own Received: header in front of the headers it got from the client. 1 SMTP transaction in, 1 SMTP transaction out. There is no Multiplication at all in this case, and this is frequently the case. Usually, they only put 10 recipients in a single SMTP transaction. The message is not sent in any less time, either. The spammers connection typically limits the speed of the run. They can send the same messages, in the same time directly to the mailserver for domain. The fact that there are many, many more users than there are domains, means that multiplication is always going to be insignificant. However, the closed relay offers exactly the same multiplication effects, however insignificant they may be. The open relay does not offer the spammer anything they don't already have.

"Recipients of the spam could then flood your server with complaints"

In practice, and in my experience, this doesn't happen. ISP mail administrators are usually able to identify the authentic Received: header inserted by our relay, identify the true source of the abuse, and make the complaint to that ISP. I have run open relays for 7 years, and very infrequently do we get complaints of this sort. All such complaints are replied to, and indicate the correct ISP to complain to. And we also block that IP address, if we haven't already noticed, and send our own complaint. Most of the time, these complaints arrive well after we have detected, blocked, and complained about the activity.

"If your server crashes from this overload, repairing it could be costly and time-consuming"

This is true as a general statement about crashes. In my experience however, the only overload's related to open relay abuse we have experienced have involved anti-spammers trying to teach us a lesson. It is true that any open relay operator has to be knowledgeable and prepared to deal with abuse, regardless of the source or purpose behind it. About 1 year ago, I was arguing with two extreme anti-spammers, and shortly thereafter, our server was bombarded by connections from 2400 different IP addresses. This lasted for about 10 days. Our monitoring software immediately blocked almost all of this, but it still involved sorting millions of messages. Through this, no legitimate messages were lost, and our servers were usually up to the load. This would have been a disaster for anyone unprepared for dealing with such abuse.

Shortly afterward, I began to experience abuse of the form where abusers would try to send out KLEZ virus infected messages, forged to have my From: address. There are several things to note:

- 1) Forges dean@av8.com in the return-path: header.

- 2) Is sent through odie.av8.com (198.3.136.136 is not MX for av8.com) See logs below. This IP address used to be an outbound address for av8.com several years ago. One blacklist in Europe still advertises this name on their web page. The virus sender either knows the name odie.av8.com, or has received an email address from me, but from some time ago. The fact that it is not the MX address for av8.com is significant. If an person outside of Av8 was looking for our mailservers, typically one would lookup the MX records for Av8.com. Such a lookup gives different address, and so we know they didn't do that.

3) The message sent contains a KLEZ virus. KLEZ sends out spam and forges email addresses, but it doesn't ordinarily both abuse my open relays, and forge my address. KLEZ usually forges addresses selected from the infected user's mailbox.

4) In a recent case, this week, a virus was sent through our relays from an ISP in San Antonio Texas. It sent to only 2 addresses. Later, it resent another message to one of the 2 previous email addresses. This is fairly unusual.

5) KLEZ also goes through the address book of the user, and tends to send a lot of email, not just a couple messages.

The 5 unusual behaviors leads one to think that the message was sent by a person, probably the (or perhaps just one of many) KLEZ virus operator. This is personally directed at me. This type of address forging is usually meant to harrass the person whose address is forged. In the old days, people would get spam and block the sender address. Most people have realized that this is a bad idea, since the From: and Return-path headers are easily forged.

Certainly this activity directed against me, could be the isolated act of a few miscreant anti-spammers. It may not be related to the primary creator or creators of KLEZ. However, it is unknown whether all KLEZ activity could be the result of such miscreants, and it is unknown to what extent or to what portion KLEZ activity accounts for the spam sent out. Two facts bear examination: Over the last year, or two, it seems that spam growth has grown exponentially, along with KLEZ activity, which has also grown exponentially over the same period, and which is now reported by some to be the largest Virus infection ever.

Certainly, this seems intent on creating the impression that I sent the unwanted Virus. But I think people aren't that easily fooled. Quite obviously, the abuser sent the virus.

We sent a complaint of relay abuse and Virus infection to the ISP responsible. We got this response:

On Wed, 14 May 2003, Sam Dibrell, Jr. (World Net Admin) wrote:

> Dean,

>

> Thank you for contacting World Net technical support with your open
> relay information. In order to resolve this issue, I will be happy to
> include your Class C netblock in our open relay blacklist.

>

> Unfortunately, due to the fact that you blatantly refuse to follow
> best practices involving SMTP server configuration, I will be unable to
> dedicate any further time or resources to assist you in resolving your
> issue.

>

In other words, this ISP (World Net of San Antonio, Tx) thinks it is OK for their users to send out viruses, and abuse our open relay services. They think that somehow, this is our fault. This person sounds just like Christopher Neill. Since this message, we have had another attempt from a different IP address that belongs to this ISP. Apparently, they really didn't bother to reign in their abusive user, or their abusive employee. It could be the case that their user is infected, and their computer is being controlled remotely via the internet. This capability of viruses is well documented. However, it would be up to World Net to detect this, and help identify the abuser(s). That is their responsibility. Given such an irresponsible response, I have to wonder if it is possible that they could have taken the step of sending the Virus. I don't know if they would dare go that far, but the suspicion seems credible. I note that breaking into a computer with a Virus is a violation of criminal provisions of the Computer Fraud and Abuse Act. We have filed a police report with the San Antonio Police Department. The matter is still under investigation.

And that is really the core problem: The ISPs, typically residential, dialup ISPs that offer very simplistic email services, think that they don't have to police their users. They think that since their email setup is trivial, that everyone's email setup is trivial; if it's not, then too bad for us. If there is anything that encourages spammers to abuse Open Relays, it is this behavior exemplified by World Net to ignore abuse of Open Relays, even when that abuse involves commission of a Federal Felony. This is a rare and unusual response to a complaint. Fortunately, the vast majority of ISP's respond responsibly and promptly to such complaints.

"For example, a spammer can use an open proxy to connect to your mail server. If your server is an open mail relay, the spammer can send loads of spam, and then disconnect - all anonymously"

Not so. Their IP address (for the open proxy) is still in the message, as explained before. As in the case above, involving World Net, we know from the Received headers where the abuser came from. It is not anonymous. However, in that case, World Net will not police their user. The open proxy may have been abused remotely, just like a virus infected machine can be controlled remotely. However the ISP that is responsible for the open proxy, can identify the source, and they can also close the open proxy. An open proxy, unlike open relays, does not appear to have any legitimate purpose. An open proxy abuser can connect directly to other mailservers, or they can connect to the open proxy's ISP's mailserver. They did not gain anything by abusing an open relay.

While not specifically listed in the FTC statement, anti-open relay crusaders also make false claims about the necessity of blocking open relays, in order to block spam. This is a false claim frequently made by commercial open relay blacklists, which misleads their customers, and often causes their customers to unwillingly fail to block spam that is identified as being spam. As explained previously, all relays (open and closed) add the IP address of the client to the message as its being sent. If the recipient mail server checks these addresses against an IP blacklist of known spam sources, then use of a relay (again

open or closed) cannot fool the spam filter. Therefore, it is completely unnecessary to block any relay, open or closed.

The practice of blocking legitimate, non-spam email is termed "revenge blocking". While a salient fact to the purchase decision, frequently the open relay blacklists do not inform their customers that they block legitimate non-spam email. Frequently, the open relay blacklists do not tell their customers that open relays are used by many ISPs and companies for legitimate purposes, and that blocking these relays will result in lost legitimate, non-spam email. Frequently, blacklists will block entire IP address blocks, affecting many users who are merely "next to" the spammer. They do this in "revenge", to try to get those customers to complain to the ISP, or to harm the ISP's business (in violation of anti-trust, tortious interference with a contract and other claims). When we receive a complaint from our customers of blocked email, we contact the blacklist user by phone. In nearly every case, they are shocked and surprised that the blacklist is blocking legitimate, non-spam email. Frequently, this results in their user discontinuing use of the blacklist, and implementation of more effective anti-spam measures. The users frequently seem to feel misled. I would like to urge the FTC to investigate such claims made by commercial open relay blacklists as being false and misleading.

Conclusion

One reason the FTC has received such slanted information on the subject is because the FTC mail server is itself using a "revenge" list, that blocks my email, and the email of others who would speak out on this topic. I think that explains why the panel contained no one who operated open relays. The panel was loaded with radical antispammers such as Julian Haight, John Levine, Steve Atkins, and Ray Everett-Church, all of whom have previously voiced strong (and factually wrong) opinions about open relays, and whom I've refuted on public lists. I think they know who I am.

I am particularly distressed that Ray Everett-Church, who is an attorney with an ethics obligation to make sure all sides are fairly heard, would have participated in such a

stacked panel. He certainly knows who I am. There are others, such John Gilmore, who wrote the GNU Tar command, who have spoken out eloquently for open relays. Mr. Everett-Church also knows who John Gilmore is. It is extremely distressing that these people would attempt to slant the record, and mislead the FTC on such salient facts.

Thank you for this opportunity to correct the public record,

Dean Anderson
President
Av8 Internet, Inc