Comment #: 36

From: Rusty Carruth
Sent: Thursday, September 30, 2004 8:15 PM
To: Authentication Summit
Subject: E-mail Authentication Summit-Comments

First, I am the 'owner' of 2 active domains, from which I directly send and receive email (I use postfix as my MTA). My DSL ISP explicitly allows servers being run by its clients, but explicitly disallows spamming. I recieve approximately 200 to 300 spams per day, every day.

My primary comment on all this is that whatever is chosen MUST NOT UNDER ANY CIRCUMSTANCES have proprietary or patent issues associated with it.

The Microsoft proposal, with its 'poison pill' against the GPL, is a good example of a VERY bad idea that must be rejected. All such attempts to 'monetize' the internet's email system by any party must be avoided. If the U.S.A. chooses to attempt to force such a system, the likelyhood is that some non-trivial percentage of the world (including those inside the USA) will NOT adopt the system, thus either making it impossible for folks on one side to send email to folks on the other, OR there will be email gateways - which will end up gatewaying spam (and completely voiding all possible benefits of the system!).

I would like to address some of the (possibly) false assumptions that I see in the notice in the federal register.

First is the assumption that it is possible to entirely replace ALL headers in an email message.

This is clearly false. One can create fake headers for all entries EXCEPT for the MTA to which one is speaking when one is trying to create said fake headers. ***IF*** all (stupid) MTAs would RECORD the doggone IP address to whom they are speaking in the header (some do, some don't), AND if the OS on which the MTA is running would disallow source routing, then you'd have a record of the IP address of the machine which connected to the (non-compromised, one hopes) MTA closest to the final recipient. Then its 'simply' a matter of tracing back till you find forged headers...

Another limiting assumption is that the ONLY way to stop spam is to identify the originator or otherwise 'authenticate' same.

There is actually a pretty amusing method which would guarantee a huge reduction in spam. Unfortunately, it was made illegal when DOS attacks were declared illegal, regardless of purpose. Simply stated, if a moderately large (greater than 10, probably) number of people set up 'spam attractor' email boxes to which no human was EVER directed to send email (but which were nevertheless presented 'all over the internet' for email harvesters to find), and which, when emailed to, would reply back with a warning that said email box is a spam black hole and not wise to email to. IT would also notify all OTHERs of the IP address (and any other info needed) of the suspected spammer.

If any (other) email spam catcher box received email after some pre-determined time (say, 1 day), OR if the email notification bounced, then all of the members of the group would immediately launch a distributed DOS on the spam machine. Said spam machine would no longer be able to emit spam. End of problem. (Sort of :-)

Another possible solution is the idea of 'graylisting', which sends a 'temporary failure' notice to 'first time' emailers. Its somewhat involved, I suggest a search (google is your friend) for 'graylisting'. Unfortunately, its not perfect, as it only requires that the spammer WAIT an hour from the initial attempt to send email and being able to actually send it. But it shows promise.

I have also used a teergrube to great effect. Again, search for 'teergrube' for info.

A few answers to specific questions:

1 - There is no 'silver bullet'. Spammers will always try to find another way to get the spam through. Slowing down delivery, and blocking delivery seem to be quite good defenses, though. That being said, I've seen an amazing reduction in spam when I was running my (VERY STICKY) teergrube. At the time I was running it, I went from around 20-50 per day to 1-2 per day.

3 - if any additional, non-'free' (free as in speech) software is required, then *this* recipient will not be installing it.

4 - Most will likely 'throw them away'. *IF* it guaranteed that there will be no false positives, this is ok. I doubt such a guarantee is achievable, but that won't stop ISPs from doing it.

7 - For it to succeed, whatever is chosen MUST be open.

8/9 - I do not believe that any of the things I've mentioned above are proprietary or encumbered by anyone's supposed IP.

20 - Only if its open, non-encumbered, 'free as in freedom', and NOT compute-intensive.

21 - a teergrube delays mail, using the continuation line, for a predetermined amount of time (I used 1 day).
a graylister delays the 'first email' from a sender for at least 45 minutes (exact time depends upon sender's retry period). Subsequent emails, sent within 31 days of the initial one, go through immediately.
22 - neither teergrube nor graylisting affect anonymity.

I would love to participate in the summit, but am unable to pay my own way, so am most likely unable to attend.

rc