Comment Number:	516761-100006
Received:	6/24/2005 7:23:45 PM
Submitted As:	CW Web Form
Organization:	Red Hat
Commenter:	David Woodhouse
Agency:	Federal Trade Commission
Rule:	Email Authentication Questionnaire
Docket ID:	To Be Added
Attachment:	why-not-spf.html

Comments:

Why you shouldn't jump on the SPF bandwagon

Abstract:

SPF is a very broken anti-forgery technique which you should *not* implement. You should neither publish SPF records, nor check SPF on incoming email.

If you use SPF, you will be causing genuine email to be rejected. There are much better alternatives which address the forgery problem without throwing the baby out with the bathwater.

OK, so what is SPF?

SPF is *Sender Policy Framework*, formerly known as *Sender Permitted From*. It makes the assumption that mail from a given email address will only ever be seen when sent directly from certain IP addresses which are associated with that address, and it gives the domain owner a way to specify which IP addresses are 'acceptable'.

So, for example, I might state that mail from the `infradead.org` domain will only be seen from hosts with hostnames ending in '`infradead.org`', from the IPv4 subnets `81.187.226.96/29` or `81.2.98.173/30`, or from the IPv6 subnet `2001:b80:10b::/48`. The idea would be that if you ever see a mail claiming to be from an address at `infradead.org` which is coming from a *different* host, then you should know that it's a forgery.

(That's about as far as the pretty graphical but mostly content-free official explanation goes, but if you take a moment to look closer or think about it, there's a little bit more to it than that...)

Unfortunately, this assumption is ***false***. You *do* see perfectly genuine mail from my domain, from machines other than mine. This happens due to mail *forwarding*. People tend to change their ISP quite often, but don't want to have to tell everyone that they've changed their email address. So they have an account elsewhere, at a vanity domain or on another computer, and they *forward* mail from that address to whichever is their current ISP, or their employer.

SRS

There is a proposed solution to this -- the SPF advocates say that the whole world should 'upgrade' to make the initial flawed assumptions on which SPF is based come true.

Their idea is that when a server forwards a mail, it shouldn't just use the original sender's address as has been done for the last couple of decades. Instead, it should munge the sender's address as it resends the mail, faking an address at the domain which is actually responsible for the forwarding. Then, if a bounce is generated, that faked address receives the bounce and the bounce needs to be forwarded on to the original sender of the original mail.

This forwarding of bounces could be easily abused if done naïvely, so the idea is that we also include a timestamp to time-limit it, and a cryptographic hash which can only be generated locally. And of course you have to include the original sender's address; both the domain and the username. So when Bill Gates sends me an email to my vanity domain, for example, and my computers forward it to my current ISP, they'd actually end up rewriting the sender address to: `SRS0=hh=tt=microsoft.com=bgates@infradead.org` instead of leaving it intact.

This rewriting scheme they call 'SRS', and they need it to be implemented *everywhere* before they can use SPF without throwing away genuine mail.

Problems with SPF

You may have managed to spot one or two potential problems with this scheme all by yourself, but some people do seem to have trouble so let's spell them out explicitly:

- **SPF is not compatible with today's Internet...**

  SRS is not common. If you publish SPF records, you are going to be asking people to throw away

genuine email which you did actually send.

If you check SPF records strictly for anything larger than a toy domain with a handful of users, then you're likely to be rejecting genuine email which is sent to you or your users.

- **...and won't be compatible with tomorrow's either.**

  There is little incentive for people to deploy SRS. It's a cumbersome and convoluted workaround for a broken assumtion.

  Those who believe SPF to be an insanely stupid idea will strongly resist, and those who are simply uninterested in SPF have little reason to want to use it. Yet SPF require these these uninterested third parties to upgrade too -- it's not just limited to those who are participating.

  People are *very* slow to deploy new ideas on the Internet, and especially with respect to email, even when those ideas are *good*. Many people haven't even managed to deploy Extended SMTP (ESMTP) yet, even.

  There will be a porcine implementation of RFC1149 before SRS is ubiquitous.

- **SPF is not an anti-spam technique.**

  Some people claim that SPF directly combats spam. It doesn't. SPF attempts to address *forgery*. In fact, a large amount of spam rates an SPF 'pass' result, because spammers have rapidly adopted SPF for themselves.

  You still need a blacklist or other kind of trust database, to tell you which domains are trustworthy and which are not. But we already *have* lots of blacklists; it's just that we list the IP address instead of the domain name, to tell you which *hosts* are trustworthy and which are not. We don't need to break email and force everyone to upgrade to some bizarre new scheme just for that.

- **SPF is easily duped.**

  Although the uptake of SRS on forwarding sites will be low, it can still be a common spammer trick. Any spamming host can do the SRS trick to 'take responsibility' for forwarded mail. If your ISP receives mail from the address I used in the above example -- '`SRS0=hh=tt=microsoft.com=bgates@infradead.org`' -- then they have no real way of telling whether it really did come from Bill Gates via my servers, or whether it's a fake. SPF would just accept the mail, depending on my domain's reputation in the trust database. It's all about how much you trust the one server which is offering the mail -- it's not an end-to-end authentication.

So SPF is really not any better than blacklisting by IP address or HELO name. Again, the breakage it imposes is not worth it.

- **SRS obfuscates useful information.**

  The original sender address is useful information, and can be lost if an intermediate host mangles the mail by using SRS.

  A sender address may be present in a blacklist, or may fail SMTP callouts. Yet an intermediate host 'takes responsibility' for the mail in question, causing it to be accepted anyway.

  Sender addresses are also used for mail filtering. The sender address is the most reliable method of filtering mailing list traffic into its own folder; using `Cc:` and/or `To:` headers has both false positives and false negatives, and using the `List-Id:` header has false positives too. Again, SRS mangles this information.

- **By using SRS, you can be vouching for spam.**

  How good are your spam filters? By rewriting the address of mail you're forwarding so that it appears to come from your own domain, you put your own reputation on the line. You could be blacklisted for mail which you *claimed* even though you didn't send it and you have no real knowledge of the original sender.

- **SPF offers a whitelist, not a blacklist.**

  In the simple case without forwarding, SPF *does* manage to give a clear indication that a mail is valid. But that's not really very useful in practice. What we need is a *blacklist* but SPF gives us only a *whitelist*. We can't safely use that to reject mail -- we can only reject mail if it's clearly invalid, but SPF can only honestly say 'valid' or 'unknown'.

Alternatives

In summary, SPF doesn't really achieve what it sets out to achieve, and it breaks far more than it's worth. The costs far outweigh the benefits.

Some SPF advocates would claim that you can't make an omelette without breaking eggs; that the spam problem has got so bad that we *have* to break things and switch to a new email system.

That's silly for two reasons. Firstly, if you're going to break the world and make people upgrade, you might as well just switch wholesale to something like Internet Mail 2000 rather than pretending it's still compatible with SMTP. SMTP has enough other problems *(like the way bounces are done, the difficulty of implementing variable preferences for different users, etc.)* that we might as well fix it all at once, **if** we find it necessary to make such an incompatible change.

Secondly, and more relevantly, there are alternative schemes which give a measure of protection against address forgery but *without* the flawed assumptions, the loss of genuine mail, and the need for the rest of the world to 'upgrade'. You would have to be mad to choose all the problems of SPF, when other schemes are far saner can and offer the same benefits:

- **Yahoo's DomainKeys, Cisco's Identified Internet Mail, META Signatures**

  These schemes (and some others) are all fairly similar in nature. They offer a cryptographic solution to the forgery problem, where a sender adds headers to the mail with a signature, and that's verified by the recipient by checking against a published key for the domain or user in question.

  Google, Yahoo and Earthlink, amongst others, are already using DomainKeys on outgoing email. AOL are reportedly going to be doing so in the early part of this year (2005).

  These solutions don't suffer from the forwarding problem; they're compatible with email as we know it today, and offer the same protection as SPF does but without the false rejections.

  There is an effort under way to merge these into a single proposal for approval by the IETF. See the MASS web page for more details.

- **Signed Envelope Sender (SES), Bounce Address Tag Validation (BATV)**

  These proposals are based on the same idea, originally inspired by the concept of using SRS on your *own* outgoing email instead of on forwarded mail.

  For example, I never send email with `MAIL FROM:`; it's always rewritten to appear to be from a time-limited SRS-encoded address.

  This means that I can instantly start rejecting bounces targeted at the raw email address, because I know they're bounces to mail I did not actually send. And this, in turn, means that anyone doing SMTP sender verification callouts will reject faked mail claiming to be from that address, because a bounce cannot be delivered.

  Since SMTP callouts are unpopular in some quarters, SES takes the concept further and provides

alternative methods for validation of reverse-paths, and also allows for a message digest to be included in the reverse-path, to prevent replay attacks in the event of an encoded reverse-path being discovered and abused by a spammer.

By implementing SES or BATV, you can instantly stop accepting bounces to mail which you didn't send.

- **Certified Server Validation (CSV)**

CSV is a very simple method for identifying servers by their `HELO` greeting instead of just by their IP address. It uses `SRV` records in DNS to specify which IP addresses may use a given hostname in their `HELO` greeting. It also offers a method of storing the reputation database in DNS.

Since SPF can only really be used for evaluating the trust level of the host which is actually sending the mail, CSV is sufficient to use as a viable alternative.

Further reading

- Considered Harmful: SPF
- SPF is harmful
- Problems with Designated Sender
- You Might Be An Anti-Spam Kook If...

"When banks start DomainKeys or S/MIME signing all outbound mail, I promise to give up SPF and Sender ID."
-- Meng Weng Wong, inventor of SPF.

David Woodhouse
Last modified: Thu Jan 13 14:09:11 GMT 2005