Comment #: 31

From: Peter H. Coffin
Sent: Thursday, September 30, 2004 2:49 AM
To: Authentication Summit
Subject: Email Authentication Summit--Comments (Matter # P044411)

The Notices published in the Federal Register (Vol. 69, No. 178 / Wednesday, September 15, 2004 pp.55632-55636) is a very good summary of the basic situation, with one very important omission. The Internet Engineering Task Force (IETF) has been considering what is essentially a similar question of the "Sender ID" portion of what the FTC is considering. The IETF working group called the MTA Authorization Records In DNS (MARID) working group, has recent concluded thier work, stating that it was impossible to continue. There are many analyses of the conclusion and what it all means, but for those seeking more information than I can summarize here, I recommend starting with the materials at

http://www.moongroup.com/index.php?option=content&task=view&id=31&Itemid=2

Quoting from the above, the following summarizes the essential problem encountered by the working group:

"Throughout the MARID process a consistently recurring issue barred the way to any possibility of consensus on the proposals being considered. This occurred due to overwhelming evidence that the proposed solutions could not be deployed universally as a standard because the various proposals were encumbered by Intellectual Property claims. These claims arose because of a patent application submitted by Microsoft and currently under consideration by the US PTO. Though Microsoft offered to license their technology under terms acceptable to the IETF their license deliberately excluded any possibility that an open source software product could use their technology. This happened because the license Microsoft offered and the licenses used to distribute open source MTA's were wholly incompatible. Many of the corporate representatives involved with the MARID Working Group simply shrugged this off but the more numerous independent members, many of whom use, support, or develop open source software MTA's could not ignore it. This caused an insurmountable disagreement where consensus proved impossible."

Please note at this point that the vast majority of email being delivered today is handled by software written those very independant developers. The "Big Four" of mail transport software (Sendmail, Postfix, Qmail, and Exim) handle nearly every SMTP-delivered mail at some point in the process from the sending user to the receiver, and all are license in terms that are sufficiently "open source-like" to be unable to ignore encumbered intellectual property.

One point that is vitally important to make clear when advising

laypeople (that is, not mail administrators) on matters involving the Sender-ID proposal under question is the different between a "Sender address" and a "From address". An email, both body and header information is data wrapped in an SMTP envelope during transmission. During a passing of mail from one server to another along the way to delivering a message, the mail servers only generally pay any attention to the information in the SMTP envelope. They don't generally look at the data. When the mail is delivered, the SMTP envelope is discarded. The user to whom the mail is delivered does not see the SMTP envelope. This has very important ramifications in that the address that the user sees in their email program DOES NOT necessarily have anything at all to do with the sender address that was on the SMTP envelope. Neither of these proposals in any way provides any assurance that the address a user sees is where the mail came from or even exists.

I will now address the numbered points of consideration from the Federal Register notice. One point I'm going to refer to frequently in my comments is that much of what's to be gained from either of the proposals under discussion is the concept I'll refer to as Everyone Must Implement (or EMI). That is, so long as there are systems that don't comply or don't yet comply with the new standards *that are not known about beforehand*, then the gain of adopting the standard does not materialize. Foreknowlege allows exceptions to be made inexpensively ("whitelists"). Without foreknowlege, legitimate emails from non-compliant SMTP mailers are utterly indistinguishable from illegitimate spam, and both users and mail administrators are faced with the choice of either accepting it all or rejecting it all.

1. It is unlikely that either proposal would result in significantly less spam recieved by consumers. Until such time as every mail system in the world complies with these standards, email transported "for hire" by ISPs to consumers cannot in any legitimacy outright reject email for its consumer customers. It can only filter it into "Possibly-Spam Folders" and the like, as the risk of customers unhappy with false-positive rejections is high to accept. Once a consumer is convinced to look at a spam message, the spammer has achived his goal of delivering the message to the consumer, and the existance of spam itself is evidence that it takes only a very small number of people willing to accept even the slimmest rates of success to keep sending spam mail. They're only encouraged to send more. This is one manifestation of EMI.

2. Both proposals require at least writing additional Internet Standards, and in order to have any measurable impact, they will have to surplant (not merely modify) existing standards.

3. Each of the standards requires at least a marginal amount of additional hardware and software to support it. Sender-ID likely requires the least, as the overhead of a few additional DNS lookups is quite small, and DNS caching techniques mitigate it even further. DomainKeys requires encryption and decryption which are inheranly expensive processes. (That's largely the whole point of encryption: to make trying every possible key to decrypt the value far more effort than doing something else is.)

4. Mail administrators must either accept all unauthenticated messages or reject them all. While accepting the message allows it to be marked to the recipient as unauthenticated, in practice this has essentially zero effect on the amount of spam sent or received. It only has an impact on the amount of spam read, and not a complete one at that. See point 1. Again, this is an EMI issue.

5. For all practical purposes, neither system is likely to mislabel any email. Either it came from an authorized MTA, or it didn't. Either it was correctly signed, or it wasn't. This does not mean, however, that all mail so labelled is not spam. A spammer with a zombie net and a known list of the legitimate mail relay on the Sender-ID list can send emails through the legitimate relay just as easily as the emails can be sent directly. This may allow an ISP to detect that a customer's machine is a zombie, but there is little or nothing the ISP can do about it short of suspending the customer's service. ISPs are understandably very reluctant to take that step because it is almost guaranteed to lose them that customer to an ISP that doesn't care about spam, rather than earn the ISP thanks for their vigilance and aid.

6. The two standards are not interoperable (supporting only one does not mean that a system that supports only the other is in any way distinguishable from a non-compliant spammer-friendly system). They are also not mutually exclusive. It's entirely feasable for a mail system to support both of the proposals and apply them simultaneously to all mail.

7. Yes. Both of the proposals must be open-standard or they simply will not be accepted by the authors of the software that handles the majority of mail worldwide. Please refer to my opening discussion at the top of my comments re: the IETF MARID working group. EMI applies here. Additionally, please note that DomainKeys depends on strong encryption being exportable world-wide. Weak encryption (less than 1024 bits) will provide essentially no benefit as every signing key cracked means a mail domain that has been rendered into the equal of an open relay, and the spammers have those zombie networks available to be massive parallel computing networks, almost perfectly suited to the task of cracking keys.

8. The Sender-ID proposal skirts around proprietary intellectual property patented by Microsoft. See the IETF discussion at the top.

9. See #8.

10. Sender-ID utterly breaks email forwarding services that do not rewrite the sender to their own domain. Mails from these services are and must be by definition not Sender-ID authenticated. Forwarding services that do rewrite the sender domain are either taking responsibility for the email not being spam (an impractical proposition) or will quickly end up on lists of Not To Be Trusted For Authentication lists, which would be in effect exactly similar to as if they'd never bothered to implement the proposal at all. DomainKeys isn't broken by forwarders.

11. Mobile users may potentially face issues, but the issues of remote email sending vanish behind solutions already largely implemented by mobile users in general: web-based email system, Virtual Private Networks (VPNs), and secure tunneling connections. I don't see this as an additonal burden on mobile users.

12. See #11, especially web use. Corporate employees (the majority of mobile users) have adequately addressed the implications of this for their own purposes.

13. The majority of mailing lists I am familiar with already rewrite the sender information of submissions to the list and are hosted on machines which would be either authorized to send mail for the domain of the sender address they are rewriting to, or will be at least authorized to relay mail through an machine authorized to send mail for the domain

hosting the mailing list server. DomainKeys would require re-signing mails with the new sender, and would therefore be asserting that the email came from the list, not that it came from a particular person.

14. Outsourced mail organizations will be affected, but routinely deal with incoming mail as well as outgoing mail, and thus process is already in place for making necessary DNS changes (for setting MX - Mail eXchanger records). Making similar changes for outgoing mail is a small amount of additional setup, not an ongoing burden, and technically feasible.

15. Users with multiple apparent responsible identies would be somewhat affected. The circumstance could be managed on a basis similar to the outcourcing arrangements as above, but for small organizations and end users, it would likely end up being impossible for all practical purposes.

16. Webmail that sent mail and recieved mail for users based on the domain that it is responsible for would be unaffected by either system.

17. Both systems scale no worse than linearly. DomainKeys is computationally expensive and scales nearly linearly. Many mail systems that are computationally bound (opposed to network or disk-bound) are already running other kinds of spam detection methods such as Beyesian filters or keyword and format checkers, and those methods wouldn't be replaced by DomainKeys. DomainKeys validates sender, not content. Therefore, mail servers at close to their computational limits must be upgraded and probably signifigantly. Mail-handling in general isn't a computationally-intensive task, but DomainKeys makes it so. Sender-ID scales better than linear as increased volumes of mail are more likely to find hits in a DNS cache, the only expensive part of the process.

18. DomainKeys implementation would likely require some upgrades to systems. Larger organizations that have better planning are more likely to require an upgrade than a smaller organization that likely has a machine handling mail that has excess computation resources anyway. Both proposals' checking is likely to be handled on the recieving mail server, not an end user machine which may not always have network access to retrieve keys. "Legitimate" email marketers would simply deduct the cost of new hardware and increase prices accordingly. Spammers would simply use more zombie machines.

19. EMI works both ways. An ISP would only face challenges providing email services without participating in authentication proposals if email from them is actually refused by other providers. If the email is accepted and ends up in users' "Possibly-Spam Folders", that's not the sending ISP's problem. So long as there are many systems that are not participating, small numbers of large providers cannout risk rejecting mail for EMI reasons.

20. EMI affects scheduling and rollouts as well. There are still mail servers out there that haven't implemented RFC1869 - SMTP Service Extensions, published in 1995, because there's no real requirement that they do so. They are compatible enough for their purposes with only the 1982 RFC822 - Standard for the format of ARPA Internet text messages.

21. None of the proposals would significantly increase transmission times or adversly affect consumers that are customers that have adequately provisioned for the implementation of either proposal.

22. Anonymous political speech is only able to be accomplished with great care in email to begin with, so neither of these proposals affect

the difficulty of making anonymous political speech. Even unauthenticated mail provides sufficient markers for an unwary user to be traced and correlated with other activities. Spammers have already learned enough about erasing those marks.

23. The US government only has to take normal antitrust safeguards into account when reviewing the proposals. Reasonable assessement of the MARID conclusions should be part of that process. Outside of that, the MTA authors will likely not implement anything that's too firmly beholden to a single organization.

24. Any scheme or system can be compromised. Neither of the systems are paricularly vulnerable, though, and careful support can encourage them to be less vulnerable. Sender-ID can be aided by ISP's placing "Egress filters" on their border routers to block spoofed packets from leaving their networks. DomainKeys can be aided by recommending strong cryptography and allowing such to be exported.

25. Neither proposal would have more than incidental effect on "phishing". What both proposals validate is generally the sender ID on the SMTP envelope, which the normal user never sees or cares about. Only DomainKeys pays any attention at all to the From: address that the user sees, and then only when the sender is not already in the message. Most users don't look closely at emails they get, any more than they look at the HTML to find out where a link in an email is REALLY going to send them.

26. Sender-ID for outbound email is a nearly painless process even for a small business or small ISP. It's basically a single administrative change to DNS records and then can be left alone. DomainKeys requires a slightly more complicated initial setup, but if the small business is hiring someone to set up their outgoing mail, then the person they are hiring hopefully has the skills to make such a setup. I see no difficulty in expecting that skills will manifest when there's a need, and probably for no increase in price to the small business. Setting up and maintaining either proposal for inbound email is rather more complex, but there is no requirement that any organzation, large or small, ISP or not, implement the inbound portion at all. The sole reason to would be for the percieved benefit of having done so, which comes down to a matter of "how much is this worth?" And that's a decision that can only be made at the organization's own level.

27. Both standards would require more or less world-wide acceptance in order to be effective. This is the EMI issue again.

28. DomainKeys would require an internationally-accepted (and exportable from the US if developed here) strong encryption standard.

29. My expectation is that the de facto standards will have only co-incidental relationships to any conclusion reached at the Email Authentication Summit. The EMI factor will hold over any other decision, and the what will likely result is a fair number of organizations publishing Sender-ID records as they are publishing SPF records now, and a much smaller number will publish DomainKeys records as well.

30. Other key things for the market to do: Establish as common practice the egress filtering discussed above on every level of ISP and transit provider. Additionally, blocking end user computers from connecting to foreign (non-ISP) machines on the email deliver port 25 should be more widely implemented. Thirdly, the ISPs should establish responsive abuse desks and actually disconnect customers that are identified as problems until such time as the problem is fixed, and refuse business from

customers that are repeatedly problems. Finally, this needs to be a world-wide effort, rather than a national one. Spam knows no national borders, and spammer-controlled zombie networks can just as effective from France or South Africa as they can be from US DSL customers.

--
"25 grams of wafers and 20 ml of wine undergo transubstantiation and become the flesh and blood of our Lord. How many Joules of heat are released by the transformation?" --Theological Physics exam, 1997