



Comments for the FTC Email Authentication Summit, P044411

September 2004

WRITTEN COMMENTS OF DR. JOHN R. LEVINE

It is my honor and privilege to submit these comments to the Federal Trade Commission for their public workshop, entitled *Email Authentication Summit*.

I am a consultant and author specializing in consumer-oriented Internet topics. I am the primary author of *The Internet for Dummies*, the world's best selling book on the Internet, which has sold over seven million copies in nine editions in over two dozen languages since 1993. I am also the co-author of numerous other books including the recent *Internet Privacy for Dummies* (2002) and *Fighting Spam for Dummies* (2004). In these books, my co-authors and I educate readers about e-mail and spam.

I chair the Anti-Spam Research Group (ASRG) of the Internet Research Task Force under the oversight of the Internet Activities Board of the Internet Society. The ASRG is a coordinating forum to coordinate research into and development of technical measures to deal with unwanted e-mail, with broad participation of industry, academia, and independent researchers. I serve on the board of the Coalition Against Unsolicited Commercial E-mail (CAUCE), the leading grass roots anti-spam advocacy organization.

I have spoken at many professional, trade, and government fora such as the 2003 Federal Trade Commission Spam Forum the *Enterprise Messaging Decisions* conference in Chicago, May 4-6, 2004, and the *E-mail Technology Conference* in San Francisco, June 16-18, 2004.

I serve on advisory boards related to consumer Internet issues at companies ranging from Orbitz, one of the big three on-line travel agencies based in Chicago, to Habeas, a small anti-spam certification startup in Palo Alto CA.

Responses to questions

Question 1.

Whether any of the proposed authentication standards (either alone or in conjunction with other existing technologies) would result in a significant decrease in the amount of spam received by consumers.

No authentication scheme on its own will stop spam, since spammers are quite able to send spam that passes any authentication scheme. The point of an authentication scheme is to reliably identify the source

of a message. If we can do that, then it becomes possible to look up the source in a “reputation service” that provides advice on whether or not to accept mail from that source. This offers some hope to reduce the amount of received spam, since mail from sources with bad reputations could be safely rejected.

Authentication should also assist in taking legal action against senders of spam that violates CAN SPAM and other laws. One of the most difficult aspects of such legal cases is tying the spam to a particular sender, and authentication should make it easier to demonstrate who did or did not send particular messages.

Questions 2 and 3.

Whether any of the proposed authentication standards would require modification of the current Internet protocols and whether any such modification would be technologically and practically feasible.

Whether any of the proposed authentication standards would function with the software and hardware currently used by senders and recipients of email and operators of sending and receiving email servers. If not, what additional software or hardware would the sender and recipient need, how much it would cost, whether it would be required or optional, and where it would be obtained.

All of the likely proposals such as Sender ID, CSV, and Domain Keys are designed as extensions to current mail standards. Although we don't have enough experience with any of them to be sure, they all appear to be feasible to implement on existing hardware by modifying existing software. There is some concern that the extra processing will require extra equipment, but it appears to me that the savings due to less spam filtering are as likely to decrease the amount of hardware needed as to increase it.

Question 4.

How operators of receiving email servers are likely to handle unauthenticated messages.

The sensible ones will treat them the same as they do now, applying various spam filtering heuristics to decide what to do with them. No doubt some people will decide to reject all mail that doesn't use their favorite authentication scheme, but poorly designed filters are hardly a new problem.

Question 5.

Whether any of the proposed authentication standards could result in email being incorrectly labeled as authenticated or unauthenticated (false negatives and false positives), and the steps that could be taken to limit such occur-

rences.

None of the proposed systems is infallible, and it's unlikely that any system could be. The likelihood of mislabeling varies considerably from one scheme to another; it's quite high for SPF and Sender-ID, moderate for Domain Keys, and quite low for CSV. One of most important reasons we need further large scale experiments with all of these proposals is to see what their real-like failure rates and failure modes are, so we can understand how well they work and what modifications might be useful to make them work better.

Question 6.

Whether the authentication standards are mutually exclusive or inter-operable. Whether any of the proposed authentication standards would integrate with any other standards. For example, if Mail Server A is using standard X, will it accept email easily from Mail Server B that is using standard Y?

All of the likely proposals are designed as extensions to SMTP and can co-exist with each other. If a sending host doesn't use a scheme that a recipient host uses, from the point of view of the recipient host, it's the same as if the sender supports no authentication scheme at all, as in question 4.

Question 7.

Whether any of the proposed authentication standards would have to be an open standard (i.e., a standard with specifications that are public).

Any scheme must be open and available for use at no cost if it is to have any hope of broad acceptance. There are dozens of different mail packages in use on the net, including many open source packages distributed at no charge and maintained partly or entirely by volunteers. These include sendmail and qmail, two of the most popular. Neither is likely to support any scheme that requires payment or a license with onerous terms.

Questions 8 and 9.

Whether any of the proposed authentication standards are proprietary and/or patented.

Whether any of the proposed authentication standards would require the use of goods or services protected by intellectual property laws.

Microsoft has filed patent applications that cover their Caller ID and Sender ID proposal and, arguably, SPF. Yahoo has filed patent applications that cover Domain Keys. If these applications are granted as patents, users will need licenses for them. Microsoft has offered a license for the Sender ID intellectual property which the open source

community has broadly rejected as too onerous and legally risky for them to accept. Yahoo intends to offer a Domain Keys license similar to other licenses that the open source community have found acceptable.

Questions 10 and 13.

How any of the proposed authentication standards would treat email forwarding services.

Whether any of the proposed authentication standards would affect the use of mailing lists.

Schemes that attempt to tie particular source domains to particular sending hosts, such as SPF and Sender-ID, fail when a forwarding service forwards a message. Both SPF and Sender-ID have proposed work-arounds, but it remains to be seen in practice how well they work. Systems such as CSV that authenticate sending hosts without regard to the domains in mail they send, and systems such as Domain Keys that authenticate messages rather than senders both have no difficulty with forwarded mail.

Nearly all authentication schemes fail when a forwarding system modifies a message, such as when free service adds an advertising footer to a message, or an e-mail discussion list manager adds the list name to the subject line of messages sent to the list. In these cases it's probably necessary to treat the forwarding service as the source of the forwarded message, and encourage recipients to use mail software that can show both the original source of the message and the forwarder as sources of a message.

Questions 11 and 12.

Whether any of the proposed authentication standards would have any implications for mobile users (e.g., users who may be using a laptop computer, an email-enabled mobile phone, or other devices, and who legitimately send email from email addresses that are not administratively connected with their home domain).

Whether any of the proposed authentication standards would have any implications for roving users (i.e., users who are obliged to use a third-party submission service when unable to connect to their own submission service).

From the point of view of authentication, mail from mobile and roving users presents roughly the same issues as forwarding, in that the mail is sent from other than the nominal source of mail. CSV and similar host authentication schemes pose no problem. Message signing schemes such as Domain Keys offer the option of delegating signing authority to the mobile or roaming users, so they can send properly signed mail regardless of where they are. Source authentication

schemes such as Sender-ID and SPF don't deal with these situations well at all.

Question 14.

Whether any of the proposed authentication standards would have any implications for outsourced email services.

Outsourced mail services are increasingly configuring themselves to be logically part of the domain of the customer on whose behalf they send mail. If for example, an outsourced provider is sending mail for `company.com`, it will often arrange for the customer to delegate it a subdomain such as `email.company.com` so the provider can send mail from the subdomain with valid authentication.

A few outsourced providers have so many customers that they don't find it practical to get subdomains of all their customers assigned. The ESPC's Project Lumos is intended in part to put some sort of source indicator in each message so that clients can distinguish mail from various clients all sent from the same provider host. I have seen no interest in the e-mail recipient community in taking on the task of deciding which of a provider's customers are sending good mail and which are sending bad, so I don't see this as an important issue. In this case, for authentication purposes the mail all comes from the outsourced provider.

Question 15.

Whether any of the proposed authentication standards would have an impact on multiple apparent responsible identities (e.g., in cases where users send email using their Internet Service Provider's SMTP network but have their primary email account elsewhere).

This is a general form of the issues raised in questions 10 through 14, so the same answers apply.

Question 16.

Whether any of the proposed authentication standards would have an impact on web-generated email.

No. For authentication purposes, web mail is very easy to handle since all of the mail originates from the web server.

Question 17.

Whether the proposed authentication standards are scalable. Whether the standards are computationally difficult such that scaling over a certain limit becomes technologically impractical. Whether the standards are monetarily expensive due to hardware and resource issues so that scaling over a certain

limit becomes impractical.

The real answer is that until we have broader experience we don't know. Most of them are built on existing technology such as SMTP and DNS that is known to scale. but they all depend on reputation services that don't exist yet except in the most primitive prototype form.

Question 18.

Identify any costs that would arise as a result of implementing any of the proposed authentication standards, and identify who most likely would bear these costs (e.g., large ISPs, small ISPs, consumers, or email marketers).

Each mail system operator would bear the cost of upgrading hardware and software to handle whatever scheme(s) they choose to support, but those costs would likely be modest, no more than routine mail upgrades they do now.

Again, my main concern is reputation services. A useful reputation service is similar in concept to a credit bureau, expressing opinions about the people or organizations it rates. Credit bureaus don't work for free, and there's no reason to expect that reputation services would either. Some of the prototype reputation services such as Ironport's Bonded Sender require senders to post a substantial bond (reportedly \$25,000) against future spam complaints. This could potentially be a significant cost to senders, particularly to senders who are not well known and wish to send large amounts of mail.

Question 19.

Whether ISPs that do not participate in an authentication regime would face any challenges providing email services. If so, what types of challenges these ISPs would face and whether these challenges would in any way prevent them from continuing to be able to provide email services.

The more widely accepted authentication becomes, the more likely recipients are to discriminate against senders who don't authenticate. Per the answer to Question 7, I expect any widely used authentication schemes to be available at nominal cost, so the bar to implementing them would be low.

Question 20.

Whether an Internet-wide authentication system could be adopted within a reasonable amount of time. Description of industry and standard setting efforts, whether there is an implementation schedule in place and, if so, the time frames of the implementation schedule.

It took about ten years for existing proprietary e-mail systems to become integrated into the e-mail system. Despite the perceived urgency of doing something about spam, it would be unwise to try to force deployment of new standards in a time scale measured in less than years. The Internet's e-mail system is enormous with many subtle interactions among the parts, and any change is likely to have unanticipated side-effects. Pushing out changes before we understand what those side-effects are and how to deal with them runs the risk of causing major disruptions in the mail system.

Question 21.

Whether any of the authentication standards would delay current email transmission times, burden current computer mechanisms, or otherwise adversely affect the ease of email use by consumers.

No. All of the likely schemes are designed to work in real-time as messages are sent or received.

Question 22.

Whether any of the proposed authentication standards would impact the ability of consumers to engage in anonymous political speech.

No. All of the likely systems work at the level of e-mail domains or hosts, not of individual mailboxes. It is currently straightforward to sign up for a mail account anonymously at providers such as Hotmail and Yahoo, and none of the likely systems change that.

Question 23.

Whether any safeguards are necessary to ensure that the adoption of an industry-wide authentication standard does not run afoul of the antitrust laws.

Probably not. The existing IETF processes appear to be adequate to provide for open participation.

Question 24.

Whether a spammer or hacker could compromise any of the proposed authentication standards by using, for example, zombie drones, spoofing of originating IP addresses, misuse of public/private key cryptography, or other means.

Any authentication scheme will certainly be attacked by spammers. We won't know what their vulnerabilities are until we have enough experimental experience to know what the attacks are.

Question 25.

Whether any of the proposed authentication systems would prevent phishing, a form of online identity theft.

Maybe. Phishing is a somewhat different problem from spam because in a phishing attack it is usually adequate to approximately rather than exactly forge the identity of the target organization. All of the proposed schemes would prevent a criminal from sending mail with a return address of, say, `citibank.com`. But it is easy to engage in phishing without using the target company's domain. Either the criminal can use a similar sounding domain such as `citibank-accounts.com`, or else format the return address in mail messages in ways that make popular mail programs display a misleading address that resembles a trusted organization.

The most effective approach I see against phishing is industry specific branded signatures. For example, for the banking industry, the FDIC might provide signed authentication certificates to member banks that include an image of the FDIC seal. Mail programs could be modified to display the seal in a branded signature, and public education programs would encourage consumers to expect the familiar brand in mail from organizations in that industry.

Question 26.

Whether the operators of small ISPs and business owners would have the technical capacity to use any of the proposed authentication standards. Whether any of the authentication standards could be reasonably implemented by smaller ISPs.

That shouldn't be a problem. Authentication software will likely be distributed as upgrades to existing mail programs. Some configuration and setup will be needed, but it's on the same order as setting up an SSL web server, something that small ISPs seem able to handle.

Question 27.

Whether any of the proposed authentication standards would have cross-border implications.

The technical aspects shouldn't present issues, but reputation services could be sticky. For example, is it possible to run a reputation system in a way that complies with EU data privacy laws?

Question 28.

Whether any of the proposed authentication standards would require an international civil cryptographic standard or other internationally adopted standard and, if so, the implications of this requirement.

All of the cryptographic proposals I've seen are based on SSL or S/MIME, both of which are already deployed around the world. Since they don't seem to have required formal adoption by governments, I don't see why any of these schemes would either.

Question 29.

Description of how the Email Authentication Summit can support industry or standard-setting efforts.

The most important thing it can do is to encourage experiments with all of the proposals at large enough scale to draw conclusions about what would happen if they were deployed across the entire Internet. Until we have significant experience with these schemes, the risk of wide-scale deployment far exceeds any likely benefit.

Question 30.

Assuming a domain-level authentication system is established in the near term, future measures that the private market should develop and implement in order to combat spam.

Authentication helps people avoid receiving spam, but it does nothing to keep spam from being sent in the first place. As the ISP industry matures, it has to realize that it resembles banks and other industries that provide valuable services in that there will always be people who will try to misuse their services. The most effective way to prevent that is to avoid providing service to miscreants in the first place. In the banking industry, for example, most banks subscribe to services that track people with a history of writing bad checks so they can avoid opening accounts for them. At this point, spammers routinely hop from ISP to ISP, and they rarely have trouble establishing new service when one ISP kicks them off for bad behavior. If they used services analogous to the bad check writer services, they could both prevent a lot of spam and save themselves the considerable expense involved in identifying and cancelling undesirable customers.

Another area that the private sector needs to address is computer security. So long as large numbers of consumers are running software that is easily compromised and turned into spam sending zombies, it will remain difficult both to prevent criminals from sending spam and to know who those criminals are. ISPs need to establish procedures for identifying compromised customer computers, for isolating those computers from the rest of the Internet, and for helping consumers to repair their computers. Even more important, ISPs and consumers must demand that software vendors provide software that is resistant to compromise. Security designs to prevent compromise and limit the damage if compromise occurs have been well known in academia and computer industry for many decades. There is no reason why they

can't be applied to consumer software, and in today's Internet, it's increasingly irresponsible not to do so.

Revision date: 2004/09/30 05:44:42