Trusted Email Open Standard

A Comprehensive Policy and Technology Proposal for Email Reform

An ePrivacy Group White Paper

By
Vincent Schiavone
David Brussin
James Koenig
Stephen Cobb
Ray Everett-Church

May, 2003

Abstract: The problem of unwanted mass email, known as spam, which currently threatens to overwhelm legitimate email traffic—the messages that people want to receive—can be solved by reforming email through widespread adoption of the Trusted Email Open Standard (TEOS). This platform-agnostic standard combines technology and policy to create the trust and accountability that is currently lacking in email. By securely identifying email senders and enabling them to make verifiable assertions about the messages they send, including participation in programs that promote best practices, the Trusted Email Open Standard provides a solution to spam. Because this standard combines proven technology with broad consensus on best practices, adoption can be rapid, with costs more than offset by savings from spam reduction. Just as important, the Trusted Email Open Standard safeguards the interests of all responsible users of email, from legitimate bulk senders and email service providers to consumers, even those individuals who wish to use email anonymously.

For more information, contact Stephen Cobb via email: sc at cobbassociates dot com

Executive Summary

As the popularity of the Internet grew, email evolved from a novelty to a necessity. Unfortunately, the continued viability of email as a useful communication medium is now jeopardized by a rising tide of unwanted email, known as spam, which threatens to wipe out the many present and future benefits of email.

Today's Technology

All email currently transmitted across the Internet is sent using an agreed-upon industry standard: the Simple Mail Transport Protocol (SMTP). At present, any server that "speaks" SMTP is able to send mail to, and receive mail from, any other server that speaks SMTP. To understand how "simple" SMTP really is, here is an example of an SMTP transaction. The text below represents the *actual* data being sent and received by email servers (the numbers and words in CAPS are the actual SMTP instructions):

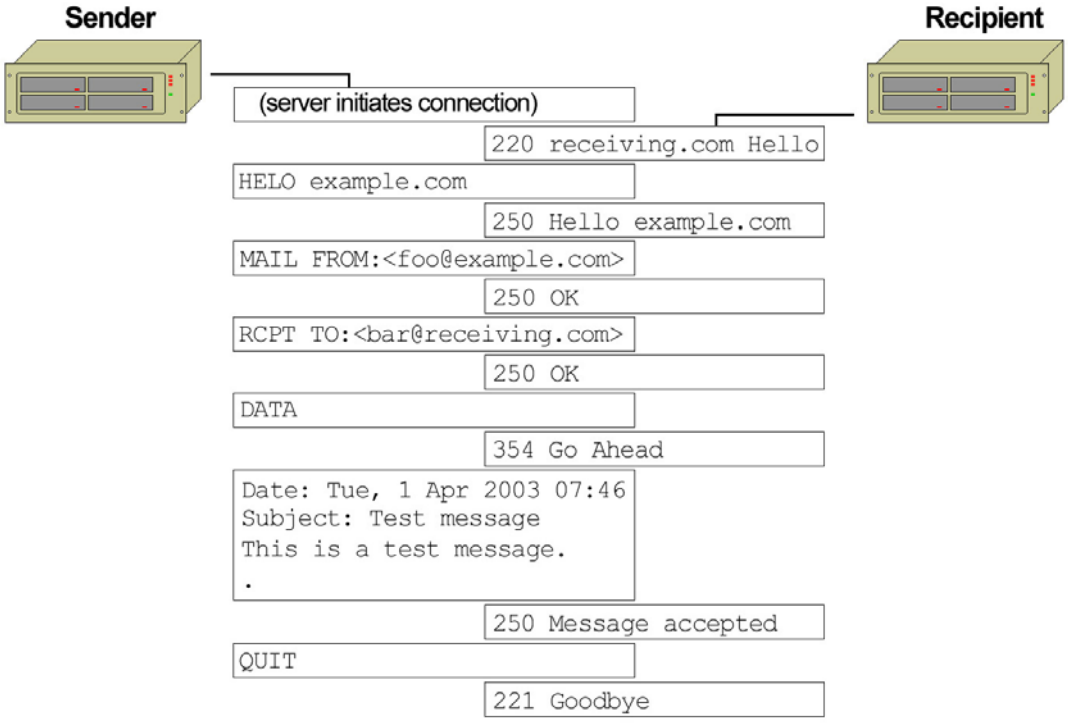

Because a server may be processing a dozen or more message connections per second, the SMTP "conversation" must be kept very brief. SMTP does this admirably, but that simplicity is both a blessing and a curse. As you can see from the example, there are only two pieces of identity information received before the mail is delivered: the identity of the sending server and the From address. Because SMTP has no process for verifying the validity of those identity assertions, both of those identifiers can be trivially falsified.

The remaining contents of the email, including the Subject and other header information, are transmitted in the data block and are not considered a meaningful part of the SMTP conversation. In other words, no SMTP mechanism exists to verify assertions such as "this message is from your bank and concerns your account" or "this message contains the tracking number for your online order" or "here is the investment newsletter that you requested."

Some sites do perform whitelist or blacklist look-ups on the IP address of the sending server during the SMTP conversation, but those inquiries can dramatically slow mail processing, requiring extra capacity to offset the loss of efficiency. Maintenance of whitelists can also be time-consuming, and blacklists have a long history of inaccuracies and legal disputes. In short, the need for speed creates a system in which there are virtually no technical consequences for misrepresentations in mail delivery. And this is precisely why spammers have been, and continue to be, incredibly effective in getting unwanted email delivered.

SMTP is, as Winston Churchill might have put it, the worst way of "doing" email, except for all the others that have been tried. The reality is that SMTP works reliably and has been widely implemented. To supplant SMTP with anything "better" means a wholesale redesign of the entire global email infrastructure, a task that few in the industry are willing to undertake. Therefore, the real challenge is to find a solution that can ride atop the existing SMTP infrastructure, allowing SMTP to continue functioning efficiently while giving those who use it the option of engaging more robust features that help differentiate legitimate mail from spam.

The purpose of this ePrivacy Group white paper is to describe a means of utilizing the existing SMTP infrastructure to communicate more information about the quality and nature of an email message, and to do so in a way that does no harm to the foundations of the SMTP infrastructure. The core of our proposal is to transmit, along with the body of the message, new information that enables senders and recipients to verifiably identify one another (something we know can be done, because we have developed such technology and it is already deployed by some companies). We believe that, until the world technology community is ready to adopt a more robust and secure mail transport protocol, this approach presents the greatest potential for realistic implementation in the foreseeable future.

Trust, Accountability, Technology

The reality of today's email system is that anybody can send as many messages as they like, to whomever they like, containing whatever they see fit to send, misrepresenting the source, content, and purpose by whatever means the sender sees fit to employ. The result is a system in which every email message is suspect and there are few meaningful consequences for those engaged in deceptive practices. In short, email today lacks two critical elements: trust and accountability.

For successful communications, the communicating parties must be able to trust in the identity of the other party (that the information is really coming from the party indicated), and trust in the content of the communications (that the information contained therein is

what the party intended to communicate). Communicating parties must also be able to rely upon the communications and hold the other party accountable for the statements made therein. In the offline world, interpersonal relationships can be the basis of trust (recognizing someone's face or the sound of their voice), as can technologies for establishing identity (driver's licenses, smart cards, biometrics).

Unfortunately, in today's email infrastructure, where forgery and fraud are commonplace, communicating parties must rely upon alternative means of establishing trust and insuring accountability. So we have proposed an infrastructure for introducing trust and accountability into email, using freely available technologies to provide methods for communicating and verifying identity and identifying content. We believe that the best way to create trust and accountability in email is to:

1. Establish enforceable standards based on best practices and compliance with applicable law.
2. Enable reliable and secure communication of sender identity and assurances about messages.
3. Create a balanced and broadly-supported Trusted Email Oversight Board to provide for ongoing oversight and arbitration of standards.

In this document, we address each of these elements as we describe the Trusted Email Open Standard (TEOS). Of course, the proposed technologies are only tools to serve a greater purpose: providing mechanisms for holding parties accountable for their actions. And for accountability to mean anything, there need to be real and significant consequences for failing to abide by acceptable standards of behavior. This has led us to conclude that without standards and oversight—for both the acceptable behavior and the enabling technologies—there can be no trust or accountability in email. Agreed principles and enforceable standards are needed at different levels, with technical measures and accountability guarantees tied to the strength and richness of the assertions being made about the messages being sent. We believe that a tiered approach is the most logical, the most likely to achieve consensus, and provides the most flexibility for adopters.

We describe three proposed tiers that provide increasing levels of assurance to recipients and greater reliability in the assertions being made by senders. We then describe how the Trusted Email Open Standard can be implemented through send and receive software components using a system of digital certificates of various types together with a range of machine-readable assertions, some of which rely upon encryption for enhanced trust and verification reliability.

We explain how this comprehensive solution to the spam problem can be achieved rapidly and at acceptable cost by means of a balanced oversight board that sets both technical and behavioral standards, and oversees a broad federation of email programs. We conclude that the trust infrastructure created by the Trusted Email Open Standard will rapidly elevate legitimate email so far above spam as to render it irrelevant, while also enabling a variety of programs designed to further enhance trust, privacy, and intelligence in email, all without interfering with the current ability of individuals to send email.

Introduction

The Trusted Email Open Standard (TEOS) combines three essential elements:

1. Best practices
2. Enabling technology
3. Oversight

Following this introduction, which outlines the problem that the standard is intended to address, three sections describe the standard, one for each of these essential elements.

The Problem

As a medium for communications, email has become extremely useful and practically universal. But email has also become a battleground. The usefulness of email today, and its potential for future growth, are jeopardized by a rising tide of unwanted email, known as spam, which threatens to wipe out the many benefits and advantages of email.[1] The pain is being felt by all parties in the email chain: those who receive email (recipients); those who send it (senders); and those involved in delivering it (providers).[2]

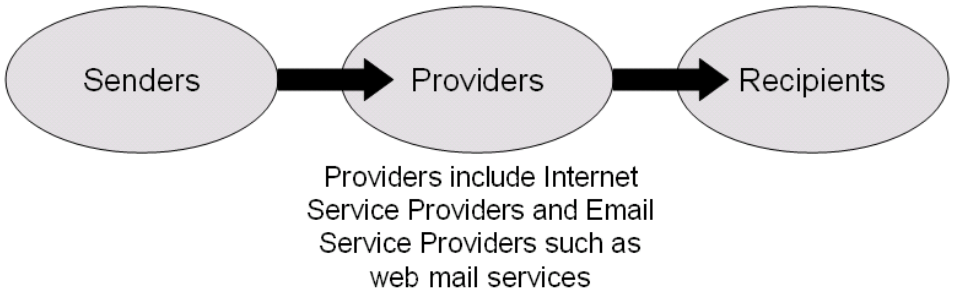

Figure 1: Simplified email chain, showing the three email constituencies

Increasingly drastic measures are being taken to deal with this problem. Recipients filter incoming messages, using add-on filtering products, email client filter capabilities, and paid filtering services. Senders of legitimate email ponder legal action to force delivery of messages in the face of increasingly frequent delivery failures due to imperfect blocking and filtering mechanisms. Providers struggle to fine-tune the blocking and filtering measures they have put in place to defend against potentially crippling spam attacks, even as they expand their capacity to cope with rising spam volumes, and struggle to deliver all the mail they should but none of the mail they shouldn't. The continuing escalation of

[1] Numerous sources indicate that spam will exceed 50 percent of all email traffic at some point in 2003 (while an agreed legal definition of spam is notoriously difficult to achieve, most people "know it when they see it" and consider spam to be any unasked for email that they are unhappy to have received).
[2] This is an intentionally broad categorization of parties to the "email chain" and will be further subdivided as appropriate. Note that recipients are sometimes referred to as email users or consumers; this category includes people using email at home or at work, for personal or business purposes; providers include Internet Service Providers (ISPs) and web mail service providers; senders include both organizations and individuals, although much of what is being proposed as a standard for senders applies mainly to companies and other organizations sending email in volume.

the "spam wars" is costly, in terms of resources and collateral damage. The latter encompasses a lot more than legitimate messages that don't get through, it also includes an increasing amount of consumer and business fraud—including personal and corporate identity theft—that employs spam as a medium precisely because it lacks accountability. Ironically, because spammers who perpetrate fraud are equally aware of the lack of trust in email, they misappropriate brand names to add a gloss of respectability to their schemes, thereby devaluing those brands and requiring victimized companies to devote resources to brand defense.

Given that spam volumes continue to increase despite anti-spam laws and lawsuits, high profile prosecution of spammers, and a plethora of technical anti-spam measures, it is clear that a new approach is needed. New laws are currently being considered, but holding email senders accountable to any principles or guidelines for conduct is currently so impractical as to be virtually impossible.[3]

This lack of enforceable standards is not a criticism of any person or entity. Indeed, the fact that email still exists today—as a fast, flexible, universal, low-cost communications medium—in the absence of enforceable standards, is a very positive reflection on the Internet community's ability to maintain a some level of adherence to self-imposed standards of conduct. However, human nature being what it is, the emergence of divergent interpretations of acceptable behavior was as inevitable in email as in all other areas of human activity. It is now time to enforce email best practices—acceptable standards of behavior—through a combination of consensus and technology.[4]

The Potential for Consensus

At ePrivacy Group we have learned that it is possible to achieve consensus on best practices, principles and standards.[5] For example, in creating the Trusted Sender program, ePrivacy Group was able to achieve consensus between commercial email senders and recipient (consumer) advocates as to what constitutes responsible email practices. A considerable number of companies have—through their participation in the Trusted Sender program—demonstrated their willingness to accept the principles and standards established for that program. However, we realize that Trusted Sender is not everyone's idea of a universal standard for responsible email behavior. In email, as in other areas of technology, the closer you get to a universal standard, the less specific the standard can be. Nevertheless, the experience of developing and deploying the Trusted Sender

[3] Despite prosecutions under state laws, and a range of successful lawsuits by providers and recipients, spam volumes are higher than ever, in terms of both absolute numbers and a percentage of total email volume.
[4] This position, and the system of trusted email oversight presented in Section 3, has recently gained backing from some significant industry players, including Microsoft. According to Brian Arbogast, corporate vice president in Microsoft's MSN and Personal Services Division, Microsoft favors independent oversight bodies to "spearhead industry best practices, and then serve as an ongoing resource for email certification and customer dispute resolution…these authorities could place a seal of approval on legitimate e-mail, making it easier for consumers and business to distinguish wanted mail from unwanted mail." April 28, 2003, (http://www.microsoft.com/security/articles/antispam.asp).
[5] The best practices supported by the Trusted Email Open Standard encompass the principles, supported by responsible commercial emailers, that are laid out in the Fair Information Practice Principles, the OECD Guidelines, and the U.S. Safe Harbor Principles.

program and its underlying Postiva technology has been instructive.[6] A number of lessons learned have been applied to the Trusted Email Open Standard.

To be successful, a solution to the spam problem must establish agreed, principle-based best practices, together with technically enforceable standards that can be applied at different levels, starting with a minimum standard. The three levels we think are most logical and likely to achieve the consensus necessary for broad and rapid adoption are:

A	A Minimum Standard for Accountability	Sender Identity (Level I) Optional Assertion of Message Type Minimal cost
B	Bulk Sender Trusted Email Certification Programs	Sender Identity (Level II) Required Assertions - Message Type - Relationship/Permission - Standardized Opt-out Optional Assertions (program dependent)
C	Consumer Oriented Trusted Email Certification Programs	Sender Identity (Level III) Required Assertions - Message Type - Relationship/Permission - Visible Assertions - Secure Seal with One-click Verify - Trusted Opt-out - Privacy Policy Link Dispute Resolution Process Trust Authority Oversight Optional Assertions (program dependent)

These three levels are described in detail in Section 1. Our technical proposal to enable low-cost implementation of this three-tiered approach is described in Section 2, while Section 3 describes a scalable system of oversight capable of spearheading broad and swift adoption of the Trusted Email Open Standard.

[6] To the best of our knowledge, no other entity has accumulated comparable experience in developing and implementing a consumer-facing, best practices-based standard for email, incorporating third-party oversight, cryptographically secure identity authentication, and real-time compliance verification.

1. Best Practices

1. Best Practices
2. Enabling Technology
3. Oversight

In this section we explain each tier of the proposed standard to demonstrate how it impacts the spam problem. Section 2 will address the technical aspects of implementing the standard. Oversight and administration of the standard will be addressed in section 3.

A. A Minimum Standard for Accountability

The most obvious gap in current email standards is the lack of any meaningful technical requirement that senders reliably identify the source of messages. The task of enforcing existing email laws and regulations is severely challenged by the current ease with which a message's domain of origin can be misrepresented or "spoofed." Bogus source and path information may be placed in the message header to prevent or undermine attempts to determine its true origin. Spammers today are using spoofing for a variety of reasons:

a. To make it harder for the recipient, the recipient's email service provider, or law enforcement, to pursue a complaint, or take action, against the sender.

b. To bypass any blacklists that have identified the sender's actual domain as a source of spam and are being used by email service providers to block messages.

c. To take advantage of any whitelists that email service providers are using to enable messages from certain senders to by-pass spam filters.

d. To defraud the recipient by leading them to think that the message is from someone from whom they want to receive messages, rather than from the person who actually sent it, and from whom they likely have no interest in receiving messages.

e. To deflect bounced message traffic resulting from their mailings, such as "Returned email" or "Undeliverable email."[7]

Because many email service providers use the most basic of anti-spam strategies—filtering out messages believed to be spam, based on a variety of factors—the detrimental effects of spoofing have required the development of increasingly complex countermeasures, notably "whitelists." These are lists of senders whose email is passed through the filtering process to avoid potentially costly collateral damage caused when

[7] Some spammers send the same message to millions of addresses at once, knowing that only a few thousand are likely to be valid, thus generating a large volume of bounce messages with which the spammer would rather not have to contend.

filters identify "false positives."[8] Increasingly, whitelists are becoming a significant drain on the resources of providers, as the implementation, management, and constant negotiation (and re-negotiation) of white-listing agreements consume the time and attention of staff.

Spoofing the source of email is also used to perpetrate a form of corporate identity theft aimed at defrauding consumers by exploiting brand names (a phenomenon experienced by Microsoft, eBay, Bank of America, Symantec, and others). Furthermore, the value of various Internet properties is diminished through fraudulent use of various domains in forged headers. For example, aol.com, ebay.com, yahoo.com, hotmail.com and msn.com are among the domains most commonly forged by spammers, resulting in wide-spread consumer confusion regarding the actual level of spam emanating from those domains.

Additionally, legitimate email service providers inadvertently add to the confusion by using common IP addresses for many clients, some of which may unwittingly use email lists of questionable quality. This often causes their email, even if it is from clients who follow best practices, to be flagged as spam by filters and ISPs. There a numerous indications that this form of collateral damage from anti-spam measures is increasing.

(i) Basic Identity Requirement: Trusted Email Domain Identity (TEDI)

Nearly a decade of spam-fighting[9] tells us that people who send spam are apt to lie about many things. Because there are laws, in the United States and many other countries, that prohibit deceptive business practices, a sizable percentage of today's spam could be classified as illegal. (The Federal Trade Commission has quoted studies that put the current figure at 70 percent.) However, acting on this fact in the absence of any verifiable data as to the source of the message, is prohibitively labor-intensive.[10]

Because the current email infrastructure creates no impediments to forgery and deception, we believe there will be little meaningful progress toward a solution to the spam problem until there is a simple, fast, scalable, verifiable means of reliably identifying the source of a message to at least the domain level. Thus, the first step towards a comprehensive solution is our recommendation that a machine-readable proof of source domain identity should be a minimum standard for email.[11]

---

[8] In the context of anti-spam filtering, a technology now widely deployed in an attempt to reduce the amount of spam that gets to the recipient's inbox, *false positives* are legitimate messages wrongly identified as spam. Because these legitimate messages are either not delivered, or at best sidetracked, they are referred to as collateral damage in the anti-spam war. Some senders of such blocked messages are inclined to sue those who falsely filter them out.

[9] ePrivacy Group executives have been involved in anti-spam activities since the mid-1990s, including playing a role in the infamous "Cantor & Siegel Green Card Lottery" incident of 1995.

[10] Numerous proposals to make truthful email sender information an explicit legal requirement have been put forward and while they would serve to make a currently deceptive practice more obviously illegal, they would remain difficult to enforce without a change to the email standards of the kind we propose.

[11] For example, the illegal use of trademarks in spam is rampant but prosecution is impeded by the lack of mechanisms within SMTP to reliably identify senders. Seeking to reduce spam by threatening to sue for trademark infringement clearly does not work. However, by giving ISPs and consumers the ability to refuse commercial email from senders who do not accurately identify themselves, as we propose, a lot less of this illegal spam will see the light of day, and any identified sender who misappropriates a trademark will be readily identifiable.

We refer to this as Trusted Email Domain Identity (TEDI) and propose that it be enabled with a Level 1 digital certificate.[12] With TEDI in place, email recipients and email providers would have a means of assessing the veracity of identity claims via email and could draw their own conclusions about email from senders whose domain identity was not verifiable. Email not bearing TEDI-compliant markers could be treated as suspect, while TEDI-compliant messages could be processed with greater confidence that they are legitimate (with a means of accountability if any are later found to be questionable).

In the Enabling Technology section below, we describe how a very low cost[13] means of implementing TEDI could be rapidly rolled out, using lightweight digital certificates and two free, open-standard software components: a Trusted Email Send Engine (TESE) that processes outbound email; and a Trusted Email Receive Engine (TERE) that processes inbound email. In the Oversight section we discuss how the TEDI standard could be policed.

(ii) Optional Message Assertions: Trusted Email Type Assertion (TETA)

The basic identity requirement—a secure, provable, machine-readable assertion regarding the source of a message—is an excellent first step towards improving email and we believe it should be established as the minimum standard to which all legitimate users of email adhere. The burden on senders would be minimal. The benefits of commercial email for small businesses would not be negated.

However, we believe that many senders of commercial email will want to leverage this minimum standard to increase the reliability of other assertions they make about the messages they send, in addition to their identity. The same infrastructure that enables assertion and verification of identity can also enable content assertions that can be machine-readable and backed by the greater accountability provided by verifiable identity.

Senders often make a variety of assertions in their messages, both in the body and the subject line. Examples include “this is information that you requested” or “you are receiving this message because you opted-in to one or more of our affiliate sites.” These assertions are common in many email advertisements; however, as many recipients have observed, the frequency with which they appear has become inversely proportional to their accuracy. As government regulators have proven, holding senders accountable for behavior that contradicts prior assertions is an effective method of curbing inappropriate conduct.[14] However, the ability to hold anyone accountable using current email technology is extremely limited, both as a practical matter (knowing against whom to file

[12] The term “Level 1 digital certificate” refers to the level of the certificate within the Trusted Email Open Standard and is different from the level designations used by some commercial Certificate Authorities (CAs). The type of certificate envisioned here as Level 1 is lightweight, with minimal proof of identity required (identity is based on existing records of registered domain ownership) and only basic revocation requirements to cover theft of certificates.

[13] Low cost is intended to mean both monetarily and computationally inexpensive.

[14] See, e.g., *In the Matter of Eli Lilly and Company* (http://www.ftc.gov/opa/2002/01/elililly.htm), in which the Federal Trade Commission brought an enforcement action for email that was inconsistent with assertions made by Eli Lilly on their web site Privacy Policy. ePrivacy Group served as technical consultants to the Federal Trade Commission in that investigation.

a legal complaint), and as an evidentiary matter once you get to court (if you cannot establish the identity of the sender).

If senders are prepared to comply with a standard for identifying themselves using TEDI, stating the nature of the messages they are sending does not present a significantly greater burden. So we have included, as an optional part of this minimum standard, the use of machine-readable assertions by the sender—placed in the message header—as to the nature of the message. An assertion would be encoded by the same open standard software component (the Trusted Email Send Engine or TESE) that processes outbound email to ensure compliance with TEDI. At the receiving ISP, an open standard software component (the Trusted Email Receive Engine or TERE) would read the assertion from the encoded headers.

Making an assertion about message type would not be a requirement in the minimum standard, but a message containing an assertion as to type would be deserving of a higher level of trust than one in which the only assertion was identity. Assertions can be very minimal indeed. For example, a sender might simply assert that the message fits one of the following categories:

Trusted Email Type Assertions	
1. Unsolicited advertisements	(ADV)[15]
2. Adult	(ADT)[16]
3. Permission-based advertisements, offers	(CRM)
4. Invoices, statements, notices and customer service correspondence	(CSC)
5. Subscriptions	(SUB)
6. Official government email	(GOV)
7. Business to business or employee	(BIZ)
8. Personal, friends and family	(FAF)
9. Non-profit, charitable	(NPE)

Note that this list merely represents possible assertions about message type that could be encoded in message headers. Neither the category names, nor their three letter acronyms, represent the proposed machine codes. Also note that the codes and acronyms are not the same thing as the subject line text that is required by some state laws (but it is important to remark that coded assertions could be matched with state law requirements to produce a form of compliance checking through the processing of coded message headers).

The benefits of including message type assertions in the message header are numerous. First, it binds the sender, whose identity is provided, to a statement about the purpose of the email. False assertions regarding the content of the message can easily be tracked back to the sender, so senders will have strong incentive to be truthful and accurate in such assertions. Again, those unwilling to make such assertions would already have little incentive to utilize TEDI, allowing recipients to treat such mail in a more circumspect fashion (some may wish to simply filter it out). The presence of both identity and

[15] This parallels an existing requirement in some North American jurisdictions.
[16] This parallels a requirement that has been proposed in some North American jurisdictions.

assertion information provides ISPs with a basic level of confidence that the sender is a legitimate entity and not an irresponsible and deceptive spammer.

A second benefit of message type assertion is that deceptive practices, such as the use of misleading subject lines to encourage the message to be viewed—for example, disguising an advertising offer as a message from a friend, name brand company, or government agency—could be more easily policed. For example, a recent marketing email sent to student loan holders by Sallie Mae was sent with a subject line "Your Response Needed," which resulted in complaints that the unsolicited message was being misrepresented as an account-related message. Had the message been subject to message type assertion requirements, the true nature of the message would have been clear.

A third benefit of message type assertion is enabling ISPs and consumers using client-side email filters to better identify:

a. those messages that recipients have stated they do not wish to receive,
b. those messages—such as statements, order shipping and tracking notifications—which should never be blocked,
c. and those messages that should be delivered only under certain circumstances (such as adult content to accounts identified as belonging to non-minors).

Our analysis indicates that when spam-filtering is based purely on "bad things," such as words and phrases believed to be indicative of spam, it is overly susceptible to "false positives" problem (the inadvertent blocking of messages that recipients want to receive, due to incorrect identification of legitimate messages as spam). By including reliable positive indicators as to message type and sender identity, filtering can be made significantly less error-prone.

Fourth, as noted above, there are legal ramifications for making deceptive assertions about the content of messages. In addition, the categories of message type could be made consistent with subject declarations which already exist in some jurisdictions. Most certainly the task of outlawing deceptive business practices would be greatly aided by requiring assertions as to message type.[17]

[17] The exact number and nature of the type assertions is obviously negotiable, however, we propose that the total number used in the type assertion standard be kept small. The assertions should track, if feasible, existing legal requirements and definitions, thereby creating the potential for automated compliance programs and safe harbor for senders who agree to be accountable and abide by the rules.

B. Bulk Sender Trusted Email Certification Programs

Widespread adoption of the standards proposed in the preceding section would greatly improve the ability of all email constituents to block, filter, reduce, or ignore spam. The level of accountability in the minimum standard would improve the level of trust between the links in the email chain (senders—providers—recipients). Indeed, if all email providers and recipients were to ignore all email that failed to follow the minimum standard, spam would already be marginalized to a significant degree.[18]

Although rapid and widespread adoption of a new standard might sound improbable, we believe that consensus on the need to act decisively is building rapidly among the larger providers and senders.[19] That is why we are providing the industry with this roadmap now, to help implement a dramatic shift in email practices. However, the minimum standard is truly that—the minimum required to make headway against the spam problem. There is much more that legitimate bulk email senders can do to improve email, beyond "merely" being honest about message source.

In fact, many bulk emailers have expressed a strong desire to do more. The standard supports these aspirations. Beyond the minimum standard, there is a standard that enables legitimate bulk senders to demonstrate their commitment to responsible email practices. Beginning where the minimum standard left off, this standard requires message type assertions, in addition to proof of identity. Beyond this, the bulk sender standard would enable email enhancement programs—Trusted Email Certification Programs—and they would require senders to make additional assertions about messages.

The nature and number of the required assertions would be established by the program's sponsor, which could be a trade association, industry group, government entity, or even a large enterprise. Senders who commit to such programs would agree to abide by their rules, subject to oversight by the program operator. The program operator could be the sponsor itself or a third-party. All program sponsors and operators would be sanctioned and supervised by the Trusted Email Oversight Board, the proposed email standards body (see Section 3).

(i) Level Two Identity

The first step to take beyond the baseline is to provide greater accountability. This can be achieved through more stringent identity requirements, over and above those proposed in the minimum standard domain name identity. This would not be expensive or burdensome and can be accomplished by additional due-diligence on the part of the issuing CA and/or third party verification of identity by the program sponsor (a role that is further defined in the following section).

---

[18] Mail from legitimate sources who had yet to implement the standard could also be marginalized, so the standard must be easy to implement on the sending side, and allow receivers to make an informed decision before discarding mail. TEOS provides for this, and for much needed positive feedback on mail quality to counter the negatives on which current filtering systems rely (e.g., filtering on suspect content).

[19] For example, "3 E-Mail Providers Join Spam Fight: AOL, Microsoft, Yahoo Seek Ways to Curtail Unwanted Solicitations," Washington Post, April 28, 2003.

(ii) Assertions in Bulk Sender Trusted Email Certification Programs

Even the limited type assertions outlined as an option under the minimum standard (1.A.ii) are of significant value in adding trust and accountability to email. A wider range of assertions about messages can even further enhance the handling of legitimate bulk email. Consider some of the facts asserted in a typical commercial email message:

- I am the originator of the message: Name of sending person/entity
- This is my email address: Email address of sending person/entity
- This is my domain: Domain of sending person/entity
- Permission to send you this message comes from: Explanation
- The list to which this message was sent came from: Source
- This message is about: Subject
- The purpose of this message is: Statement
- To opt-out of future mailings do this: Action
- To contact the sender of this message do this: Action
- To read about the privacy policies of the sender: Action

All of these statements are common in messages such as statement or payment notifications from a credit card company, order confirmations or shipping notifications pursuant to an online purchase, or "special offers" from an online merchant.[20] The need for some of these assertions arises from established practices in the world of commercial email, practices which any realistic standard must accommodate if it aspires to widespread adoption. Consider the following diagram.

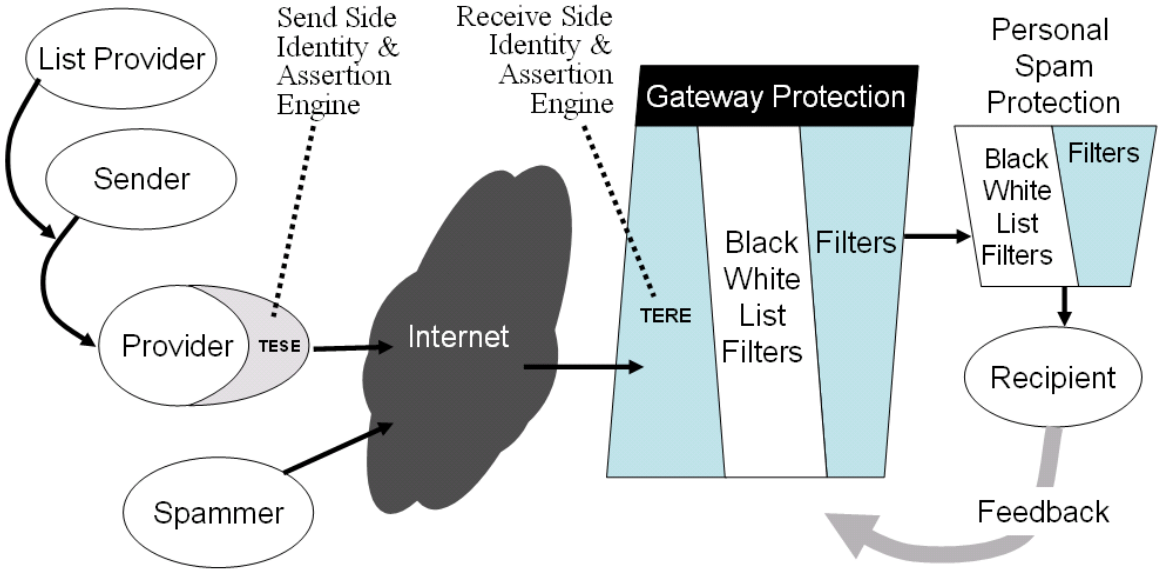

Figure 2: Detailed diagram of the email chain, showing role of list providers, and location of Trusted Email Send and Receive engines (TESE and TERE).

[20] Millions of such messages are sent and received successfully each day, to the mutual benefit of sender and receiver; but accidental blocking of them is increasing as email providers struggle to protect their resources and customers from abuse by spammers.

Information about the source of lists and their permission basis plays a vital role in the commercial email process. The ability to make, and check, assertions about this information, and other important aspects of a message, would add considerably to the trust and accountability that is currently lacking in email. [21]

The Trusted Email Open Standard would encourage and enable the creation of trusted email programs, consisting of required assertions about messages, in the form of machine-readable statements in headers, as described earlier in 1.A.i. These assertions would be in addition to placing in the message header a machine-readable assertion as to the type of message. The nature and number of the required assertions would be established by the program sponsor. As stated earlier, the sponsor could be a trade association, industry group, government entity, or even a large enterprise.[22] For example, a trade association might want its members to abide by a requirement to state list source or permission basis, so the association would establish and operate—by themselves or through outsourcing—a program for members. We recommend that responsible email senders consider the following possible assertions, based on the Fair Information Practice Principles:[23]

Notice

Was the mailing executed by the sender or an agent of the sender?
Who owns the list to which the messages were sent?
What is the permission basis of the list?
To what extent is the list shared ?

Choice:

Is a standardized means of opting-out of future mailings provided?

Access:

What access to recipient data does the sender support?

Security:

What level of protection is provided to recipient data?

Dispute:

What sort of dispute resolution does the sender support?

---

[21] See State of New York v. Monsterhut in which an email marketing firm was enjoined from sending emails erroneously claiming that they were sent pursuant to a prior "opt-in" by the recipient (details available at (http://www.oag.state.ny.us/internet/litigation/monster_hut.pdf).

[22] We anticipate that some large organizations will want to operate and self-certify their own program of email assurances.

[23] See, *Privacy Online: A Report to Congress,* Federal Trade Commission, 1998 (http://www.ftc.gov/reports/privacy3/fairinfo.htm)

Multiple assertions would assist email providers and recipients in determining the relative desirability of accepting, delivering, or reading such messages. The proposed standard would not include any technical means of verifying the validity of the assertion, only a means of making the assertion in a machine-readable way. However, the assertion would still be of considerable value because the domain level identity of the sender is—according to the recommended Level 2 standard—provable. For example, if the assertion was made that a recipient could opt-out of future mailings, and the recipient found that the opt-out did not work, identifying the responsible party for purposes of further contact or dispute resolution would not be difficult. The assertions under Notice could assist a recipient who feels they have received a message in error to determine the cause/source of the error. The Notice assertions would also be useful to legitimate senders as a means of tracking the performance of list and mailing service providers.

(iii) Standardized Opt-out

The inability to opt-out of unwanted mailings is a major source of current consumer frustration with email and a significant part of the spam problem. A recent FTC study noted that 63% of spam they received had non working opt-out mechanisms. While many privacy advocates and anti-spam campaigners have long held either opt-in or confirmed opt-in to be the only acceptable premise for commercial mailings, there is evidence that many consumers are less concerned about full opt-in than reliable opt-out. In other words, if they could reliably opt-out of future mailings they might find some forms of unsolicited commercial email acceptable.

At the same time, many direct marketing companies have advocated effective opt-out as an industry standard. In addition, ISPs and email filter companies have called for standardization of opt-out mechanisms used by legitimate bulk email senders to enable recipients to quickly and reliably opt-out of any unwanted email. Therefore, we find that all constituencies, including the Federal Trade Commission, are agreed that simple and reliable opt-out should be provided in all commercial email.

We recommend that provision of a standardized opt-out be required by all Bulk Sender Trusted Email Certification Programs. A supporting requirement would be the assertion, in the message header, that the message includes standardized opt-out. The presence of this assertion would not only aid consumers using client-side filters but would increase the ability of ISPs and filter companies to differentiate legitimate bulk email from spam and would greatly improve delivery of legitimate email.

C. Consumer Oriented Trusted Email Certification Programs

The programs enabled by the standards described in the preceding sections 1.A and 1.B would do a great deal to distinguish legitimate email from spam and enhance everyone's ability to manage email more efficiently. Spam would be increasingly marginalized.[24] However, some senders want to go even further than those standards. They would like to display visible proof to consumers of their adherence to a set of privacy principles and email best practices.

Such proof can be provided by placing a cryptographically protected seal or trust stamp in the message. The technology to do this is already in use. Several commercial senders are participating in a third-party program that provides recipients with visible evidence of their commitment to higher standards, much like privacy seal programs on web sites.[25] Tests indicate that programs of this type more than pay for themselves because trust stamps in email create a significant increase in consumer confidence, which is reflected in higher message open and respond rates, with fewer opt-outs (see the Appendix to this white paper for actual test results).

The Trusted Email Open Standard therefore includes Consumer Oriented Trusted Email Certification Programs that enable these higher goals to be met by those who aspire to them. We propose an open, accessible, non-proprietary means to achieve this. Because these aspirations have the potential to elevate email to a highly trusted status truly worthy of the term Trusted Email, an essential component of this standard is a visible means of displaying and verifying such claims.

In addition to third-party seal programs, we expect that some senders, particularly those who feel that their own brand is sufficiently trusted, will want to provide their own stamp or seal as evidence of compliance with higher standards (such stamps also provide a valuable layer of protection against corporate identity theft). The trust infrastructure of the Trusted Email Open Standard provides for such self-certifying programs.

(i) Level Three Identity

Moving beyond levels 1.A and 1.B means even greater accountability. We propose secure accountability be provided through stringent identity requirements in the form of a fully-verified digital certificate to be used by senders of email at this level.

---

[24] The economic basis of spam—the desire of the spammer to make money— means that it is subject to the law of diminishing returns. Less spam being delivered means lower rates of return. While the initial effect is for spammers to send even more spam in the hopes of getting more messages delivered, there is bound to be a point—during the decline in spam delivery rates that TEOS is intended to induce—at which this strategy becomes futile and a significant number of spammers seek alternative means of making money.

[25] For example, Microsoft MSN enrolled in the Trusted Sender program some time ago and has successfully incorporated trust stamps in tens of thousands of emails sent to customers. AES, one of the largest financial aid organizations in America has now sent well over a million trust stamped messages to its customers.

(ii) Verification: Visible, Real-time, Interactive

Those who participate in Consumer Oriented Trusted Email Certification Programs will provide visible and verifiable evidence of their assertions to recipients. These assertions could be encoded, and the cryptographically protected image (seal) that is the visual indicator of program participation could be generated, by an enhanced version of the Trusted Email Send Engine (TESE) that processes outbound email to ensure compliance with TEDI. This standard does not dictate what assertions senders should make, however senders are subject to Trusted Email Oversight Board oversight. The associated technical standard enables any assertions that are made to be verified, in real-time, through an interactive process.

Verification can be performed by the recipient and optionally by email providers. For example, an email user who receives a message from a sender who is participating in a trust program at this level may be presented with a clickable trust seal in a message. This seal enables real-time, interactive verification of the sender's identity and compliance status. An email provider may choose to verify messages as they arrive and present the recipient with a trust indicator, such as a custom inbox icon,[26] indicating that the veracity of the sender's claims has already been established. Receiving ISPs can use the Trusted Email Receive Engine (TERE) to read assertions from encoded headers and verify seals.

The practical details of how this verification standard can be implemented are explained in the Enabling Technology section of this document, but the basic operation involves two steps:

a. Placing cryptographically-encoded, machine-readable statements in messages as they are sent, typically in the form of a consumer-visible, legally-protected image or seal that conveys the nature of the assertions made.[27] For example, "This is Superior Email from XYZ," accompanied by the trademarked XYZ logo, ©XYZ, and so on.[28]

b. Providing a return communication channel that enables the decoding/verification of the statements as well as a direct connection to the sender for
    a. an explanation of the assertions it is making,
    b. opt-out from future mailings,
    c. complaint process.

(iii) Principles for Consumer-Facing Trust Programs

To be successful, any attempt to solve the spam problem through the setting of standards must remain flexible with respect to the principles it attempts to uphold. The more specifically you attempt to embody principles in standards, the less widely the standard

[26] Email service provider Mailshell was the first to implement this ability, using ePrivacy Group technology.
[27] Note that the use of an image, which has inherent protection through copyright, plus a trademarked logo/name/phrase, extends the protections against, and potential remedies for, spoofing of authorized email, into well-established areas of law, both national and international.
[28] In this example, "Superior Email" could be a program sponsored by a third-party, or self-administered program operated by XYZ company.

will be accepted. That is why the standards we have described in 1.A and 1.B are focused on one principle: accountability. This is expressed as domain level proof of sender identity plus one or more machine-readable assertions as the nature of the messages sent. Senders are free to embrace additional principles through further assertions, but they are not required to do so.

Similarly, when we get to a higher standard of email, in which senders make visible and instantly verifiable assertions, either independently or as a participant in a third-party program, we do not think that, in general, the standard should mandate what those assertions are. Of course, there are some fairly obvious candidates for inclusion by any organization that is interested in establishing consumer trust:

> **a. Accountability**: although accountability has already been addressed to a certain extent by the time you get to this level of the standard, additional accountability assertions can be made, such as "XYZ agrees to protect the privacy of customer data," or "XYZ promises to respond to recipient complaints within 24 hours," and so on.
>
> **b. Privacy Policies and Practices**: some senders may wish to signify their commitment to a set of privacy policies and practices (of their own creation or created by a third-party).
>
> **c. Third-party Oversight**: some senders may see value in submitting to third-party oversight various aspects of their email activity, such as address sourcing, privacy policies and practices, customer responsiveness, and so on.
>
> **d. Guaranteed Opt-out**: all consumer-facing programs for responsible email must guarantee that messages include a simple and reliable opt-out mechanism.
>
> **e. Dispute Handling**: senders must provide recipients/consumers with a formal process for handling complaints about messages they receive; this allows the sender to respond to, and correct, unforeseen problems arising from mailings.[29] Some elements of dispute handling can be automated but all programs at this level must provide for eventual human involvement in disputes that are escalated beyond an automated exception-handling process. Some non-profit, third-party entities, such as TRUSTe, have extensive real-world experience resolving consumer disputes in the privacy space and could be contracted to provide this service.

This list is not exhaustive, but we propose Accountability, Guaranteed Opt-out, and Dispute Handling be required by this standard.[30] To summarize this section we have graphed the 3 tiers of the proposed standard on the following page.

[29] Consumer dispute resolution is different from arbitration of disputes between email constituencies with respect to standards, which is to be handled by the Trusted Email Oversight Board (see Section 3).
[30] We also suggest that consideration be given to participation in a global "Do Not Email" list which could be established as part of this standard and potentially avoid the nightmare scenario of state-by-state DNE lists.

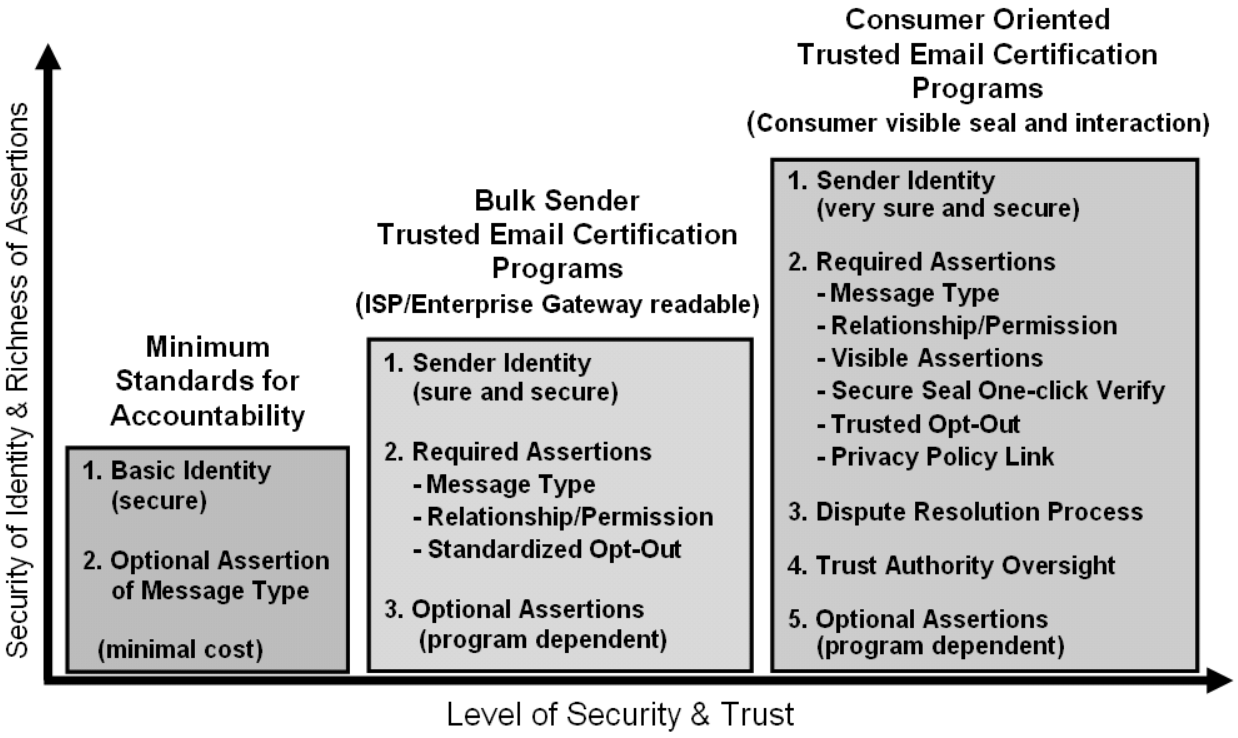

Figure 3: Overview of the 3 levels in the Trusted Sender Open Standard.

2. Enabling Technology

1. Best Practices
2. Enabling Technology
3. Oversight

Technology plays a vital role in ensuring that standards of behavior are met, principles are adhered to, and regulations are complied with. This is already true for some aspects of the Internet today. For example, if you try to sidestep the requirements of the domain name system and create a rogue domain, it does not work. Unfortunately, the current set of email protocols and technical standards, cannot effectively enforce compliance with what most people consider to be minimum acceptable standards of behavior. Email technology must be extended and enhanced at a very basic level. In order for this to happen, technology enhancements must be low in cost and easy to implement. That is why we are making elements of our proprietary technology freely available (for example, the Trusted Email Send Engine or TESE, and the Trusted Email Receive Engine or TERE).

Requirements

Practical implementation of the 'best practices' and 'oversight' components of the Trusted Email standard requires the satisfaction of a number of technical requirements. This section describes these requirements and provides a level of implementation detail necessary for the definition of a technical standard.

Just as the requirements for the non-technical components of the standard are broken down into three sections (A Minimum Standard for Accountability, Bulk Sender Trusted Email Certification Programs, and Consumer Oriented Trusted Email Certification Programs), this section will address the technical elements as they relate to this structure.

A. A Minimum Standard for Accountability

This baseline section of the standard requires the communication of certain data by senders/sources of email messages to gateway receivers of those messages. It creates several technical requirements related to this data communication.

1. Securely verifiable statement of sender identity

Sender identity requires that the apparent SMTP source/sender of a given Trusted Email message be verifiably the same as the actual source/sender of the message. Sender identity further requires that receiving systems be capable of detecting fraudulent Trusted Email messages, and messages from Trusted Email compliant domains which are not Trusted Email compliant.

Given the structure, implementation, and common practice in the SMTP protocol and other Internet communications, this sender identity will identify the domain source of

email messages, and will utilize the existing domain ownership and registration hierarchy to aid implementation.

It is recognized that the technical source of email messages is often different from the communicating entity, as is the case with commercial messages delivered by service providers on behalf of corporations. The sender identity defined by this standard recognizes this sender/source distinction and requires that both the sender and source of each message be identifiable.

2. Revocation of identity

Sender identity must be revocable in the event of an identity failure. Identity failure is defined as erroneous, fraudulent, or otherwise inappropriate association of a sender/source with an Internet domain.

3. Sender-generated declarative statements

Each compliant message must contain certain sender-generated declarative statements in order to communicate the data required. Compliant messages may also contain additional declarative statements, both sender-generated and otherwise.

Declarative statements must minimally include the capability to express statements of the following complexity: [parameter] [logical operator] [value]

Declarative statements must be constructed using elements defined in the Trusted Email standard or extensions thereto, in a structure defined therein. Additionally, declarative statements must express information within schemas or structures defined in the Trusted Email standard or extensions.

Declarative statements must be easily and efficiently able to be processed by receiving and intermediary systems.

4. Non-replayable trust elements

All trust elements satisfying the above requirements must be constructed in such a way as to preclude replay attacks. For example, a given message from Alice to Charlie should be readily identifiable as fraudulent given use of trust elements taken from a message from Alice to Bob.

5. Non-repudiation of trust elements

Compliance with the principles of the Trusted Email standard requires accountability of sender/source actions, and hence requires non-repudiation of all sender-generated trust elements and of sender/source identity.

B. Bulk Sender Trusted Email Certification Programs

In addition to those defined above, this section of the standard adds requirements related to communication of additional data.

1. Securely verifiable 3rd party trusted identity

Extending the minimum standard identity concept, the addition of 3rd party trust requires that a sender/source be verifiably the same as an entity identified by a certifying 3rd party. Compliant messages at this level might have multiple 3rd party trusted identity statements. Receivers of compliant messages may individually verify trusted identity statements as appropriate for processing.

2. Revocation of 3rd party trusted identity

Unlike the more basic minimum identity, individual 3rd party trusted identity statements may be revoked by their certifiers. Criteria for revocation will vary by certifier, and may include issues other than basic identity failure.

3. Securely verifiable 3rd party trusted declarative statements

An extension of the sender-generated declarative statement, these statements are made by a 3rd party in reference to a securely identified sender/source as described above. The intent of trusted declarative statements is to encapsulate statements by certifying 3rd parties such that they can be depended on by receiving and intermediary systems. The core construction and structure of these statements is as defined for sender-generated statements.

4. Revocation of 3rd party trusted declarative statements

As with trusted identity statements, trusted declarative statements may be revoked by their certifiers. A requirement for discrete per-statement revocation within a body of statements made by a given certifier has *not* been identified at this time, thus implementations may force 'all or nothing' revocation of statements by a given certifier.

5. Special case: Self-certification

A special case of the 3rd party trusted identity and declarative statements described above has been identified for 'self-certifying' entities. Such entities, including governments and certain commercial entities, will leverage the strength of their existing trust relationships and not require additional 3rd party certification. Technically, this special case will not differ from nominal use of the standards defined for this section, as the requisite 3rd party will be defined as a self-certifying extension of the appropriate sender/source.

C. Consumer Oriented Trusted Email Certification Programs

This level of the standard requires the communication of certain data to individual recipients of email messages in addition to gateway receivers.

1. Display of trust mark to recipient

The display of a trust mark is a representation to the recipient of message that some set of verification and processing actions have been successfully completed on an appropriate set of 3rd party trusted identity and declarative statements

2. Special case: Self-certification

Once again, the self-certified case does not require special technical consideration other than the satisfaction of technical policy requirements for this program level.

Standardization Overview

Given the requirements described above, this section will address components of a technical standard satisfying those requirements.

A. Basic Components

1. Cryptographic Medium

Cryptographic protection of data, particularly though the use of asymmetric cryptography, provides an elegant medium for the satisfaction of multiple requirements expressed above.

Key header:
Base64(<publicExponent>):
Base64(<modulus>):
Base64(signed_digest(Base64(<publicExponent>):
Base64(<modulus>)))

Signature header:
Base64(signed_digest(<senderAddress><rcptAddress>
<Assertions>)<message_specific_data>)

2. Cryptographic Algorithms

RSA Public Key Cryptography
SHA-1 Hashing

Message generation requires one SHA-1 hash operation and one RSA signature operation per message. Message verification requires a minimum of one SHA-1 hash operation and one RSA signature verification given cached pre-verification of signing keys.

3. Performance Considerations

A standard implementation achieves approximately 130-150 messages generated per second on a single 1Ghz desktop class x86 CPU. The same CPU can produce approximately 1600-2000 message verifications per second.

Since most MTA's are I/O bound and have excess CPU capacity, cryptographic verification can outperform DNS-based verifications in many architectures. Of course, DNS verification is available as an optimization to simplify adoption and accelerate high volume receivers.

4. Declarative statement language

Structured with a logically equivalent XML-based human-readable representation and compact machine-readable representation. A namespace-based standard structure and extensibility provide the technical framework for a durable and federated standard.
For example:

Human and Machine Readable Form
Namespace declarations
<assertion>
<parameter namespace> message_type </parameter>
<operator namespace> is_equal_to </operator>
<value namespace> customer_service </value>
</assertion>

Bytecode Form
AC040001AA4828C1

How the Standard Works

This section describes technical details of the Trusted Email Open Standard in key four areas: distributed verification, identity confirmation, program participation/standing, exception handling. Although each is described separately, all four work together to provide a solid basis for marginalizing spam and solving the spam problem.

Technical Considerations in TEDI and TESE

The domain name and identity forging in spam that is one of the most dangerous and prevalent elements of the spam problem stem in large part from the lack of any organizational hierarchy of identity for the email transport. Other transports on the Internet rely on one organizational hierarchy of identity: the ICANN-administered domain/DNS infrastructure. While DNS is certainly used for email delivery, the complex asynchronous nature of SMTP means that the domain/DNS infrastructure does not present the barriers to abuse that are more evident in synchronous protocols (HTTP, etc).

In fact, security exists today on the Internet largely in the form of SSL, which leverages the one-to-one, synchronous nature of HTTP connections and the domain/DNS infrastructure to add cryptographic transport protections. This suggests that a workable and relatively secure solution to these forging and spoofing problems can be made available—in the form of a generalized, foundational trust framework for email—by cryptographically linking the Internet's existing organizational hierarchy of domain names to the email transport through an application of trusted email technology. This is what refer to as the Trusted Email Domain Identity (TEDI) framework.

TEDI addresses a key weakness in the email transport: domains would, for the first time, have a means to prevent fraudulent use of their information in spam messages.
In addition, spam filtering solutions would benefit tremendously from the availability of trusted header information in email messages, as current systems make delivery and

blocking decisions based on spoofable IP address information, resulting in the blocking of innocent senders and the exploitation of white lists to enable delivery of spam. Content analysis techniques (Bayesian, rule-based, pattern matching and otherwise) could also be applied to greater benefit with the advantage of trusted domain and path information from email headers, since probability and score-based systems would naturally be more accurate given trusted data elements.

The relative merits of the TEDI approach can be seen from the following table, which compares TEDI's header-based, signed assertions with two alternative approaches. The first is DNS/IP-based sender identification. The second is the implementation of S/MIME for identity assurance:

	DNS/IP	S/MIME	Header-based Signed Assertions
Weight and Message Size	Negligible	Very High	Low
Computational Expense	Very Low	Very High	Low
Security	Low	Very High	Very High
Per-Message Assertions	Not Possible	Possible	Yes
Persistence	Not Possible	Yes	Yes

The TEDI framework would leverage existing CAs, existing domain/DNS registries and registrars, and the proposed Trusted Email Oversight Board infrastructure. CAs would issue X.509v3 certificates to domain owners through the organizational hierarchy of the registry-registrar-domain structure. These certificates, which we refer to as "level one" certificates, would have a "trust level" that is different from certificates currently issued by CAs, as they make no assertion about identity other than legitimate status as domain owner through the domain/DNS hierarchy.

Based on these "level one" certificates and their Trusted Email Oversight Board signatures, standards compliant email components, which could be cheaply implemented using the Trusted Email Send Engine (TESE) software to which we have already referred, would encode the trusted domain identity information in outbound email. The Trusted Email Receive Engine (TERE) software would read the trusted domain identity information from the headers of inbound email. With unique, cryptographically secure headers for each individual message, the originating domain would be identified to intermediary and recipient systems.

Note that a mechanism for including individual-level signatures to augment domain identity also exists and could be deployed as the Trusted Email Individual Identity (TEII) extension. This extension to the standard could provide the benefits of domain identity at a more granular level. Of course, this level of granularity would require the use of individual-level trust infrastructures, such as Passport and Liberty, to link individual email senders' identities to sending systems.

While it is possible, as we have demonstrated internally, to utilize fully standard-compliant X.509v3 certificates with our technology, the compressed signature hierarchy that we are proposing—eliminating the recursive verifications typical to certificate-based systems, including for individual identity—together with lightweight signatures, permits real-time processing, unlike other X.509-based systems. For example, un-optimized benchmarks show verification performance of around 2060 email verifications per second on a 1Ghz Windows XP desktop system. TEDI thus has the potential to more efficiently eliminate the damage done by forged headers in spam, making fraudulent messages significantly more difficult to perpetrate.

Implementation of the TEDI framework would have a broad and immediate impact, including an increased ability to fight spam on many fronts. The cost barrier to implementation is also quite low. The required "level one" certificates have few of the issuance costs associated with traditional certificates, as identity verification is not required. Certificate Authorities and registry/registrar companies may even decide to issue these certificates free of charge for an initial period—perhaps one or two years—as part of their contribution to the fight against spam, realizing that this standard would lead to a vast new market for certificates, significant renewal revenue, and new sales of higher-level certificates to domain owners.[31]

Distributed Verification and TERE

The Trusted Email Receive Engine (TERE) software allows email providers to translate, and verify the source of, machine-readable information placed in messages by senders. That information consists of source domain identity and a variety of assertions, such as the type of message, its permission basis, and so on, as described in sections 1.A and 1.B. We are prepared to make this software freely available to Internet Service Providers in order to implement this standard. The software accomplishes the following as it processes messages:

a. Authenticates the identity of the message sender by reference to a domain name certificate.
b. Reads assertions made about the message by the sender.
c. Checks for the sender's participation, and real-time standing, in any email trust program under which the message was sent.

As a result of the above, the technology enables the ISP to accomplish the following:

a. Verify message compliance with any applicable laws;
b. Determine sender's compliance with ISP terms of service (which may include requirements as to message type, accurate header information, and so on); and,
c. Enforce such recipient/customer-specified parameters as the ISP may elect to support.

[31] The Consumer Oriented Trusted Email Certification Programs level of the standard would benefit from TEDI as well, since TEDI compliant domains would need no new technical components, only a new signature on their certificate following Trusted Email Oversight Board certification, in order to send email under a Consumer Oriented Trusted Email Certification Program.

Consumer Visible Seals

We have proposed (1.C) that organizations establish consumer-oriented email trust programs under the auspices of an independent oversight body. These programs would require senders to make certain assertions as to message type, content, purpose, origin, and so on. A participant in such a program, or an entity that has been authorized to conduct its own program could, for example, install an appliance[32] configured with the requisite software, in between the participant's network and the Internet.

The installed appliance would automatically create an individually encoded Trust Stamp or seal for every designated email message sent out by the program participant's email server. These stamps would be placed in the messages as visible evidence of the participant's agreement to comply with the specific requirements of the program, such as email best practices and a dispute-resolution process. Note that this technology, which can reliably determine whether or not email senders ascribe to a set of guidelines, already exists and is not something that needs to be developed in the future. In current implementations for HTML-based email, created with ePrivacy Group's Postiva software, when someone reads a stamped email that they have received, the Trust Stamp is visible in the top right corner as an image. The image, which was generated uniquely by the TESE, contains the email address of the sender and recipient, along with the date the message was sent, and the trust mark that identifies the program.[33]

If the information appearing in the Trust Stamp matches that displayed by the recipient's email program, there is some level of assurance that the message is genuine. However, the recipient can always click on the Trust Stamp to activate a complete verification process via the official program web site. The recipient does not have to use any special software or plug-in to do this and nothing is downloaded to the recipient's computer.

When a recipient clicks on the Trust Stamp, an interactive form at the official program web site compares the information in the encoded stamp with that found in the message itself. This includes the To and From addresses as well as the Subject of the email. Based on the recipient's responses, and an automated exception handling process, the message is either verified as authentic, or verification is denied and the recipient informed of the basis for this denial (with notice of the denial also going to the program operator for further investigation). Note that the TERE technology discussed above includes the ability for ISPs to process stamped email and flag the message upon delivery to the recipient with an icon or other indicator that the message has a certain status (for example, "Customer Service Correspondence").

Exception Handling

ePrivacy Group has demonstrated that email identity verification technology works, and that real-time verification of program standing works. This has been demonstrated on a large scale, in multiple real world deployments. In doing so, it has been possible to verify

[32] By using an appliance, the participant is not required to perform any hardware or software configuration, easing implementation for less sophisticated participants. For more sophisticated installations, Postiva software could be integrated into existing systems.
[33] If messages are text-based, or viewed as text, the user is presented with a clickable text link.

the expectation that some people will attempt to abuse such technology, and some events will confuse the technology. And of course, despite everyone's best efforts, it is perhaps inevitable that the technology will confuse some people. Therefore, a means of dealing with these issues—collectively referred to as "exceptions"—will be needed. Given the sheer scale of email volumes, an automated means of dealing with exceptions is essential.

To this end ePrivacy Group has developed technology that is able to determine, through automated processing of message recipient input, the nature of the exception. Furthermore, based on the nature of the exception, this technology is able to resolve or escalate. There are two basic categories of escalation: technical and dispute. A technical escalation occurs when either a significant technical problem or a serious attempt to abuse the system is detected. A dispute escalation occurs when a significant difference of opinion between sender and receiver is detected. Both technical and dispute escalations require a human response on the part of the trusted email program operator in the case of a technical escalation, on the part of the program's designated dispute resolution operator in the case of a dispute escalation.

3. Oversight

1. Best Practices
2. Enabling Technology
3. Oversight

Whenever there are principles or standards to be met, or rules and regulations to be complied with, some form of oversight and enforcement is required. Traditionally, oversight is accomplished by entrusting the function to a body created or designed for this purpose. Examples include the Securities and Exchange Commission, the Insurance Institute for Highway Safety, the American Arbitration Association, the International Telecommunication Union, and numerous others.

These bodies require a considerable degree of independence from the entities they oversee and must evoke trust in those for whom they perform the oversight, and in those whom they oversee. Due to the pervasive nature of email, any attempt to enforce principles and standards for email requires a trusted oversight body that has widespread support, encompassing all relevant interests, from recipients (consumers), to email providers (ISPs and web mail providers), to email senders (companies, government agencies, non-profits, and so on).

Without input from all constituents, the trust body will have difficulty establishing its independence, its trust, and its authority. This section describes the broad-based trust body envisioned by the Trusted Email Open Standard: namely, the Trusted Email Oversight Board.[34]

A. Trusted Email Oversight Board Role

We propose that the Trusted Email Oversight Board function at 3 levels, corresponding to the levels of standard described in the preceding section. This is diagrammed at the conclusion of this section.

(i) A Minimum Standard for Accountability

Role: Root identity, set and evolve policy and technology standards, arbitrate disputes

At a basic level, the Trusted Email Oversight Board has responsibility for setting standards. This includes determining the principles that guide behavior that are embodied in, and implemented by, the technical standards.

[34] In this context, it is interesting to note recent comments by Brian Arbogast, corporate vice president in Microsoft's MSN and Personal Services Division, and the executive sponsor for the privacy pillar of the company's Trustworthy Computing initiative: “While we believe any related [email trust] technology implementation should be based on open standards, the creation of this independent email trust authority would be a significant step in the right direction.”

Given the broad range of constituents who have a stake in email, disputes are likely to arise in the implementation of standards and the Trusted Email Oversight Board will serve as the arbitration body that resolves such disputes (note that this role is very different from the handling consumer disputes with respect to assertions about email or other email practices—we are not proposing that consumer dispute resolution be handled by the Trusted Email Oversight Board).

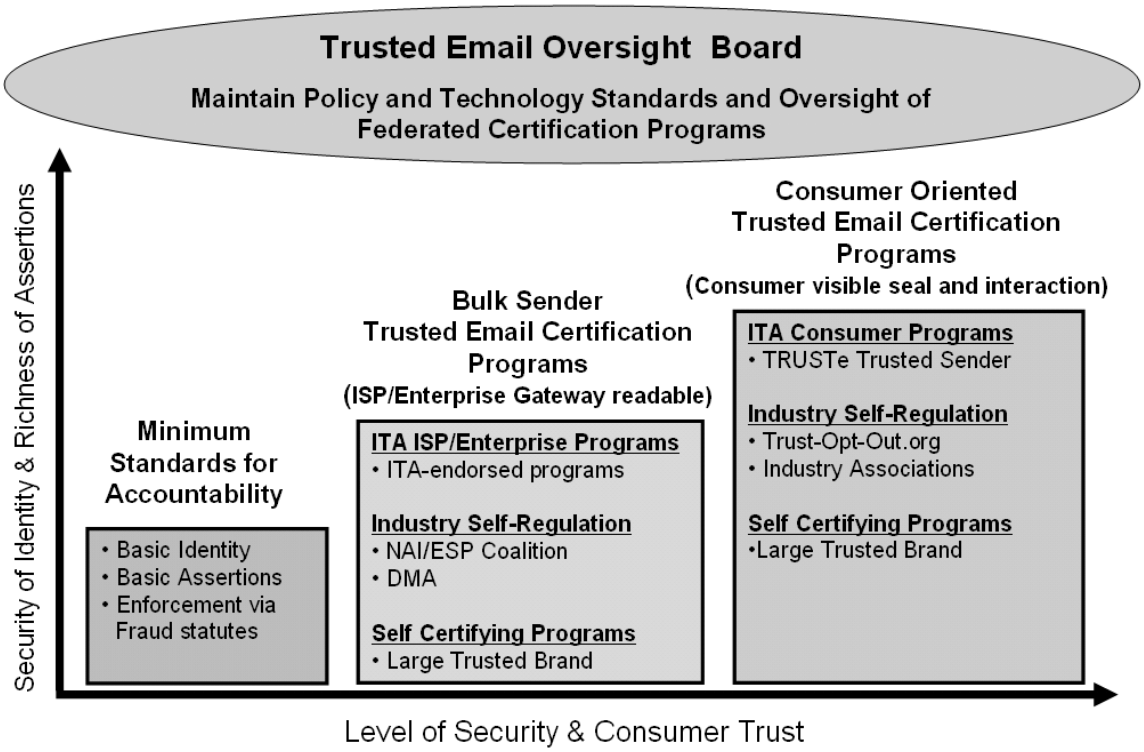

Figure 4: Relationship of the Trusted Email Oversight Board to the 3 tiers of the standard.

(ii) Bulk Sender Trusted Email Certification Programs

Role: Root identity, program sponsor and operator authorization, standards enforcement, oversight

We expect that many organizations will want to go beyond the minimum standard described in 1.A and develop programs at the bulk sender level defined in 1.B. The Trusted Email Oversight Board will have the task of approving such programs. The program sponsor, the organization wishing to create and run the program, would apply to the Trusted Email Oversight Board for approval.

While the standards embodied by the programs would be determined by the program sponsor, the Trusted Email Oversight Board would have the authority to determine the

fitness of the program sponsor to run the program.[35] The Trusted Email Oversight Board would also approve program operators with whom sponsors may contract to operate programs on their behalf.

In the event of unresolved or egregious complaints from ISPs or other constituents, about email sent under a particular program, the Trusted Email Oversight Board would arbitrate and enforce a resolution, up to, and including, revocation of program approval.

(iii) Consumer Oriented Trusted Email Certification Programs

Role: Oversight of verification, program sponsor and operator authorization, dispute resolution

When organizations wish to create Consumer Oriented Trusted Email Certification Programs and offer visible verification of their commitment to consumer privacy and email best practices, the Trusted Email Oversight Board will authorize third parties to implement, operate, and oversee programs that provide this level of trust and assurance. These third parties will undertake consumer dispute resolution with the Trusted Email Oversight Board only becoming involved in extraordinary circumstances.

B. Trusted Email Oversight Board Structure

We propose the following design for the trust body that will oversee the Trusted Email Open Standard, the goal being: credibility, balance of interests, and ability to scale, including international expansion. This structure is diagrammed on the following page.

(i) Broad-based and Board-based

We propose that the Trusted Email Oversight Board operate on a board-basis, with a broad membership, to include consumer groups, privacy advocates, and industry representatives. This structure is flexible, scalable, and readily adaptable to international operations through adjustments to board membership.

(ii) Delegated Authority

The Trusted Email Oversight Board will authorize other entities to engage in activities supportive of the standards, such as sponsoring and operating email trust programs, issuing certificates, and so on. Authority would flow from the Trusted Email Oversight Board to these entities, with the Trusted Email Oversight Board performing an arbitration role in the event of disputes between entities.

[35] The program sponsor may elect to run the program itself or outsource operation of the program to an approved program operator.

(iii) Third Party Operation

The operation of the Trusted Email Oversight Board with respect to implementation of authorizations, decisions, programs, and so forth, could easily be contracted or outsourced, possibly through an industry-funded coalition of the willing.

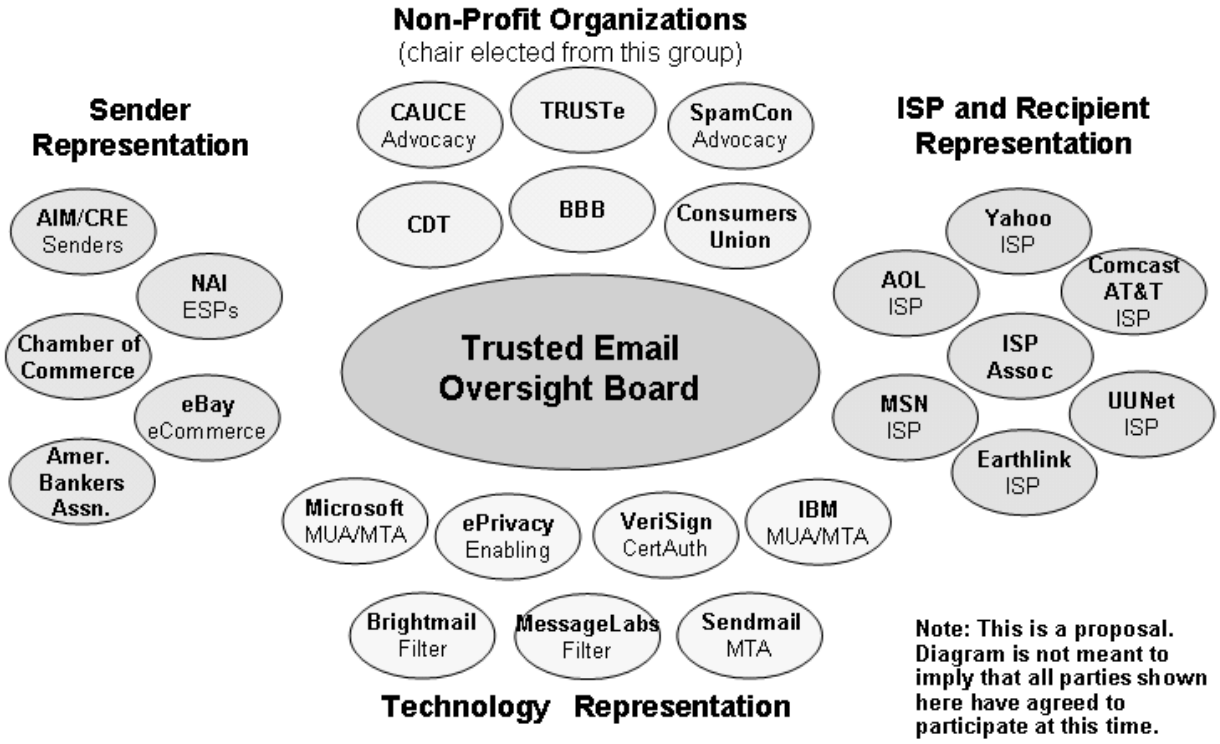

Figure 5: Proposed structure of the Trusted Email Oversight Board.

Conclusions

Without prompt and concerted action, the pain of the current spam problem will only get worse. Spam volumes will continue to climb, imposing increasingly severe financial burdens on ISPs and enterprises. The anti-spam arms race will escalate. Filters and false positives will continue to plague legitimate emailers and frustrate consumers. Legislators will try to outlaw spam even as bulk mailers seek laws that force ISPs to deliver their email. The growth of email—which has so far been the Internet's "killer app"—may falter, as spam-afflicted consumers give up in disgust.

Fortunately, a comprehensive solution to the spam problem can be achieved. We think this white paper describes that solution: the Trusted Email Open Standard (TEOS). Based on direct experience developing and implementing email technology that enables verification of sender identity and assertions, and distilling many person-years of experience—including detailed discussions with industry veterans, some of the largest players in the industry, as well as the leading consumer advocates—TEOS is, in our opinion, the best way to create trust and accountability in email because it:

- Establishes enforceable standards based on best practices and compliance with applicable law.
- Enables reliable and secure communication of sender identity and assurances about messages.
- Creates a balanced and broadly-supported Trusted Email Oversight Board to provide for ongoing oversight and arbitration of standards.

For well over a year now we have been advocating, both publicly and privately, the basic principles of this approach. We have gained the support of a wide range of consumer advocates. We have seen a steady growth in support for our approach from marketers and providers (most recently from three of the world's largest ISPs: AOL, Microsoft, and Yahoo). However, until the industry takes concrete steps to implement new standards, they will remain an exercise in wishful thinking.

We think the key is a broadly-federated system of email trust programs, guided by a Trusted Email Oversight Board that can rapidly implement a framework of technical and behavioral standards built on freely available enabling technology. This will elevate legitimate email so far above spam that spam will be rendered irrelevant. At the same time, a variety of programs will be enabled to further enhance trust, privacy, and intelligence in email.

Appendix: Field Tests of Trusted Sender

Trusted Sender is an example of the kind of program that is possible within the third level of the Trusted Email Open Standard, known as Consumer Oriented Trusted Email Certification Programs, described in 1.C and 2.C of this white paper. A field test of the Trusted Sender program, which places visible trust seals in outbound email, was conducted in the first quarter of 2003 by a large and well-known consumer company (hereafter referred to as Consumer Company).

The test consisted of emailing a consumer offer to two groups of 20,000 customers. The messages sent to the control group did not contain the Trusted Sender seal, whereas those that were emailed to the test group did.

Economic ROI

The seal had an overwhelmingly positive impact on the mailing. Compared to the control group, the test group had:

- 23% higher open rate[36]
- 52% higher click-through rate per delivered email
- 61% lower opt-out rate per delivered email

These numbers show that the higher overall click-through rate of 56% for the test group was impacted due to people both opening the mail at a higher rate, and then also responding at a higher rate once they opened the email. The test group was also much less likely to opt-out of receiving future offers. During the test, only 2 recipients opted out. Zero complaints were received and no disputes reported.

Trust ROI

A week after the offer was mailed, both the test and control groups were mailed a survey. Invitations to take the survey were sent to a total of 38,800 customers, of whom 2,631 responded (just over 9%). The responses indicated that the seal's impact on trust was also overwhelmingly positive.

- Over 80% said that use of the seal would definitely or somewhat increase their ability to differentiate legitimate Consumer Company email from spam.
- Some 79% said the seal would definitely or somewhat increase their comfort-level that emails from Consumer Company are truly from Consumer Company.
- 76% said the seal would definitely or somewhat increase their level of trust that Consumer Company respects their communication preferences.

[36] All results significant at a 99% confidence level except for click through rate per viewed email, significant at an 80% confidence level.