Message Verification Platform

Respecting Provable Relationships and Imputing the Personal Value of Interruption
A Point-By-Point Response to Arguments Against a Pure Economic Solution

Philip Raymond, Co-Chair
Email Accountability Initiative
Marlborough MA
508-864-4800
phil@senderatrisk.org

Marshall Van Alstyne, Co-Chair
Email Accountability Initiative
Cambridge MA
508-864-4800
marshall@senderatrisk.org

Our presentation format is modeled on a Q&A between a skeptic and a proponent for a pure economic solution to Spam. Questions have been adopted from statements, articles or panel debates with John Levine, Richard Clayton and Dave Crocker.

ABSTRACT

A 2006 announcement by AOL and Yahoo offering preferential treatment of paid email [GoodMail model][1] raise legitimate concerns about the potential emergence of an economic underclass. Is the era of nearly free email coming to an end?

Existing economic solutions to spam are polluted by non–free-market elements in an attempt to overcome perceived shortcomings. They cost senders of desirable mail, rely on private trust registries and a community definition of spam, and they pay the wrong parties. More complex economic models substitute risk or refunds and arbitrate prior relationships, but they are perceived as complex or offer little benefit until widely adopted.

The ideal implementation of an economic mechanism respects existing relationships, prevents collection scams and zombies, and automatically adjusts to accommodate the *personal value* of interruption. **Risk** becomes a pure expression of intent, while **Trust** and the value of interruption become market commodities rather than the award of a centralized registry. Senders who correctly assess the desires of their recipients are never taxed.

The mechanism will be modeled in a hypothetical case study and defended by addressing skeptics point-by-point. The case study demonstrates an implementation that is virtually transparent to senders and recipients, uses existing SMTP protocol, and provides benefit before widespread adoption. A final barrier to deployment was eliminated this month with adoption by the Email Accountability Initiative, a collaborative of inventors and programmers creating a unified specification and the tools, test platform and services to exploit a standardized implementation.

1. INTRODUCTION

Professor Marshall Van Alstyne presented his *Attention Bond Mechanism* to Microsoft Corporation in late 2003. One month later, Bill Gates announced to the World Economic Forum that a sender risk model would facilitate the end of spam within two years. Although the prediction was premature, the economics are solid and central to a development project that thwarts spam while delivering all desirable mail – even automated and commercial mail.

In 2008, the project is being adopted by the Email Accountability Initiative and transformed into an open-source project co-chaired by Prof Van Alstyne and Vanquish CEO, Phil Raymond. The sender-risk model uses a transparent market-feedback loop to empower recipients while preventing abuse and overcoming adoption issues.

2. HYPOTHESIS

The idea of using economics to control spam raises the specter of an electronic stamp or other transport cost – at least for some senders. In February, AOL and Yahoo announced the adoption of such a postage system to assure delivery of commercial mail from pre-qualified senders. [1]

Sender pay models are easily dismissed as an unfair taxation on an inherently free medium. They cost senders of desirable mail, rely on community policies and private trust registries, and they pay the wrong parties. More complex models substitute risk or refunds and arbitrate existing relationships, but they are perceived as complex or – at the very least – appear to offer little benefit if not widely adopted.

A pure economic model can be fair, simple and totally effective. This presentation describes a perfect model and responds to questions that apply to any economic mechanism.

3. SKEPTICS RAISE THESE ISSUES

3.1 From the recipient perspective:

Our Spam Conference presentation format at will be modeled on a Q&A between a skeptic and a proponent:

- What about 'Caller ID?' Isn't that sufficient?
- Who decides what is spam?
- Who keeps the money?
- Can I exempt existing relationships without complexity? Can I revoke the exemption?
- Can the model accommodate the unique needs of high-profile or high-net-worth recipients?

3.2 From the sender perspective:

- What commercial mailer would agree to these onerous terms?
- Why must I risk payment to exchange mail with individuals who know and trust me?
- What stops an unethical recipient from luring and collecting cash from an innocent sender?
- My list is clean. How can I analyze the potential cost benefit of meeting the risk-demand of every recipient on a large mailing list?
- What about desirable contact from senders who cannot overcome the financial barrier?
- This model seems to offer an offense (assured delivery) rather than a defense (antispam).
- Does the mechanism prevent 'zombies' form misappropriating equipment and cash accounts? Who pays then?
- Can it be effective before widespread adoption?

These questions have been resolved in the Message Verification Platform, a system that is virtually transparent to senders and recipients, adds very little overhead to existing SMTP servers and benefits all parties – even before widespread adoption. The final barrier to deployment was shattered with adoption this month by the Email Accountability Initiative.

4. SPAM: A Universal Definition

To thwart spam, we must first agree on a definition, and then determine what facilitates abusive behavior. Few people agree on a definition, but many associate it with mail that is unsolicited, commercial or sent in bulk. Yet everyone acknowledges that desirable mail often exhibits one or more of these characteristics. The real bane of email is not that a message is commercial, sent by a stranger, or even that it was sent by someone without a verifiable reputation. A more basic definition can lead to an effective method of dealing with it:

> *Spam is any message that you wish you had not received*

That is, spam is any contact that is irritating, harassing or simply irrelevant on a personal level. It is characterized by the fact that an individual recipient finds it to be personally objectionable. This simple definition leads to a very simple logical construct:

- Spam is undesirable mail
- Undesirable mail is a product of poor targeting
- Poor targeting is rewarded by economic incentive
- Solution: Create an economic disincentive

5. SENDER ID: Not the Real Problem

Proponents of sender identification claim that spam will dry up if senders are identified or at least forced to use a genuine and traceable email address. The phone company uses such a "Caller ID" method today. They routinely intercept messages that lack Caller ID for recipients who choose to screen calls. But the phone number of an unrecognized caller says nothing about the relevance and timeliness of the message content to the needs of the recipient. The call could be from a relative in a hospital or from a marketer with no better demographic data than a phone book. The only reason we are not swamped by thousands of untargeted phone calls each day is because of the cost and effort associated with each call. [2]

6. RISK MODEL: Sender Accountability

Suppose, instead, an intercept message said this:

> *"Your Caller ID is not recognized by the party you have dialed. If you complete the call and the recipient finds your contact undesirable, they may press *77. This will add a $2 fee to your phone bill."*

In the above scenario, *77 triggers an "interrupt fee". In effect, it says *"I found your message to be irritating, harassing or irrelevant, and so I will prod you to either refine your address list or deliver better content."* In a two-year trial of voluntary sender liability applied to e-mail between strangers, recipients rarely used their power to penalize senders. Instead, de facto filtering occurred in the mind (and in the pocketbook) of the individual senders.

6.1 Hypothetical Case Study:

The Perfect Mailing List

A jeweler files for bankruptcy and the court erases their debts. 10,000 independent retailers and individual consumers around the world lose thousands of dollars each in pre-paid deposits.

Two years later, a foreign company purchases the name and mailing list of the bankrupt jeweler without assuming their debt. As a public relations maneuver, they offer to fulfill the original order of every creditor at no charge, or alternatively, return their original deposits without obligation.

If the new company reaches out to individuals with whom they have had no prior contact and no opt-in consent, does it constitute spam? Perhaps. An ideal mechanism leaves that determination to individual recipients, and compensates each error in sender judgment at terms tailored to the individual recipients.

The company sends email to everyone on their list, but upon analysis, they learn that 80% of the mail was lost to filters and challenges. These were not seen by intended recipients. The company budgets $5,000 for print advertising hoping to reach at least the commercial buyers. But the audience is international and includes many who are not in the trade. Magazine ads could never target recipients as well as the complete list of creditors that they already have.

The company decides to risk $100 on a mailing sent through an ISP or service that has added the *Message Verification Platform tool* to their mailing system. The sender can also add the platform to his own PC or email server.

6.2 Message Verification Platform

6.2.1 The platform processes the mailing list before sending. It determines how many end users are "platform aware." In this case, it determines that 2,000 recipients use email services that will route the message around spam filters if the risk expressed within each message meets the calculated interrupt value of individual recipients.*

6.2.2 The platform reports to the sender that a budget of $100 is insufficient to guarantee delivery to all empowered recipients. To guarantee delivery to every protected recipient, it would require $750. But the report includes this advice:

…1 Maximum risk vs Likely cost:

If the email campaign pleases recipients as much as a prior campaign, the total cost is estimated to be $23. $20 represents payments to ISPs for testing the higher cash risk expressed within messages, and $3 is the predicted cash "seizure" by two recipients who will deem the contact irritating or irrelevant.

…2 If the sender wishes to minimize exposure, he can begin with a subset of the email list. After 3 days, the warranty expires and the sender can risk the same money again. Additionally, recipients who reply to the mail or click hyperlinks without penalizing the sender form a trust bond. The sender is no longer at risk when writing to these recipients unless this trust is revoked.

…3 Very few recipients (fewer than 0.5%) have seized sender guarantees at a rate that exceeds 1-in-50. [3] If the sender excludes these recipients and also any recipient with a vanity interrupt value – that is, one that was not calculated by the mechanism in real time – then the entire mailing can be sustained by a risk of $650 and a probable net cost of only $20 (no penalties).

* 6.3 How is Interrupt Value calculated?

In a fluid market, the value of any popular product or service floats with supply and demand. Likewise your attention to a stranger – no matter how brief – has value. It is a commodity that can be bid upon.

Why would recipients wish to be irritated by the highest bidder? Wouldn't that maximize profit rather than minimize spam?

With the correct input and formula, an automated bidding feedback mechanism passes only those messages that you personally would have preferred (the ones that are currently diverted by spam filters).

The amount needed to compensate a recipient for wasting his time is a product of three factors:

6.3.1 The number of senders seeking to contact a given recipient over a fixed unit of time.

6.3.2 The senders' respective confidence in their knowledge of each recipient's preferences.

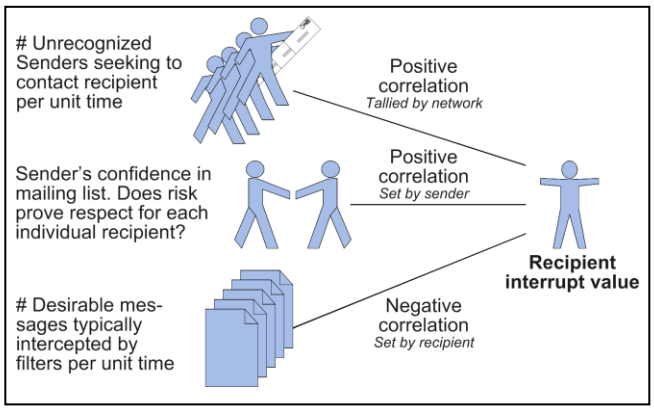

Personal Interrupt Value is a Product of Three Factors

6.3.3 Each recipient's maximum acceptable level for allowing unfamiliar contact.

This last item is the only data that a recipient needs to consider. It is simply the number of desirable messages that are typically trapped by the email provider's spam filter each day or week.

6.4 Why Filter? Is verifying risk sufficient?

Spam filters are an integral part of an economic system.

The Message Verification Platform is a sender-guaranteed delivery system. It facilitates relevant contact by waiving personally desirable messages past traditional obstacles. The delivery of desirable mail and the blocking of undesirable mail are two sides of the same problem…

6.4.1 The sender guarantees relevance directly to each recipient

6.4.2 The provider tests and routes guaranteed message past filters (typically at the recipient gateway).

6.4.3 The Message Verification Platform creates is an attention bid feedback loop. This makes it very likely that the most desirable messages will be the ones that present the minimum obstacle (decision process) for senders who have good information about and respect for recipients.

6.4.4 The recipient input sets the cutoff to match the number of messages that would otherwise fail traditional filters or block lists.

6.5 Respecting Provable Relationships —Overlooked in CertifiedMail

When a sender issues mail to a recipient with whom he has no provable two-way relationship, a certificate backed token is automatically embedded within the message to each risk-aware gateway. The token does not identify the sender. Rather it identifies and confirms one of two things:

- the level of risk expressed by a sender – and the fact that it is backed by sufficient cash at the time of receipt – or –
- the existence of a provable, prior two-way relationship

Senders never pay for guaranteed delivery of mail to someone with whom they had a two way relationship – unless the recipient has revoked the free pass.

The conversion from risk to personal trust is transparent and builds rapidly. Because the risk of 'reaching out' to someone new is limited to three days, and because engaged recipients tend to respond to desirable contact, a typical user – one who knows enough about his mailing list – requires very little money to create even massive mailings. They simply pace their contact with strangers over a short time. Larger mailing houses will put up funds for large mailings that include new addresses, but the lien will expire quickly.

6.6 Is This Layered Onto SMTP? —What Is The Resource Overhead?

The basic risk mechanism of the Message Verification Platform is already used at a dozen ISPs and hundreds of enterprises. One company in Massachusetts offers both personal software and server appliances. [7]

Risk is conveyed within a standard one-way email message and with very little overhead – typically less than 2% of CPU and data resources. When trust is established (either the parties have exchanged mail with each other, or the recipient has knowingly opted in at the web site of the sender), computational and communications overhead is further reduced. Trust relationships do not require communications with any other servers.

7. THE SKEPTIC'S QUESTIONS

7.1 What about 'Caller ID? Is it sufficient?

Provable sender identification is useful only for existing contacts (friends and family). Since spammers rarely know the address of your friends, an automated white list typically provides the same value as provable identity.

But white lists simply defend a gated community. Traditional antispam methods cannot faithfully discern desirable contact form unrecognized senders. Everyone needs contact from unrecognized senders, because it is the basis of inquiry in business, a free press and even social relationships.

Like sender identification, the Message Verification Platform uses encrypted tokens. But instead of identifying the sender, they identify risk or *pre-existing trust*. Thus, the preservation of anonymity is an ancillary benefit of the mechanism.

7.2 Who decides it's spam?

The recipient. He or she is the only one who knows what content is relevant to current needs and priorities. The recipient's decision is instantaneous, irrefutable and irrevocable.

7.3 How can the recipient decide if he has not yet seen the message?

The Message Verification Platform is a system of deterrence not penalties. Senders know when they have sufficient information about everyone in their mailing to warrant pushing their mail past filters. Risk that floats with the personal value of each recipient keeps them honest. When they are uncertain, they won't risk their wallets. If they simply have poor judgment, they will quickly run out of money.

7.4 Who gets the money?

The recipient. His or her attention is a commodity and the personal interruption value is it's fair value. Interrupting a recipient requires compensating him with either personally relevant material or cash.

7.5 What stops an unethical recipient from luring senders just to collect cash?

Even an ideal filter can only be perfect at blocking. It retards poorly targeted content by cannot reward personally relevant content. The elegance of an pure economic solution is that it provides positive benefit by *promoting* desirable contact.

If a recipient attempts to seize cash after receiving relevant contact or simply adopts a new address frequently, senders are discouraged from making contact no matter how relevant to the recipient's individual needs. Recipients who take money from the table quickly find that no one places money on the table – even in response to recipients who publicly seek information.

This suggests that recipients might be frightened to penalize truly irritating contact. In practice, relevant contact is censored at the sender. Legitimate senders know that honest recipients have a very low collection ratio.

7.6 What commercial mailer would agree to these onerous terms?

The only unsolicited commercial mailers that would agree to these terms are the ones with good information about the new addresses on their mailing list. They are very certain that each recipient will desire the content of their message.

7.7 Must I risk payment to exchange mail with individuals who know and trust me?

The answer relies on the concept of "exchange" and "trust". Senders never risk money for guaranteed delivery of mail to someone with whom they have a prior, two-way relationship unless the recipient has revoked their free pass. For marketers, this means that retention mail gets delivered without payment or even risk. Their only exposure is in the acquisition of new customers. If they seek assured delivery, they had better have good information about these prospects and respect their individual needs and preferences.

7.8 What is 'bid-for-attention'? Do senders really bid?

Not consciously. A bid-for-attention mechanism is a an economic feedback process in which the personal interrupt value of each recipient is automatically calculated based on the quantity and confidence of unrecognized senders seeking to reach him. The recipient specifies a number of permissible interruptions (typically, the number of desirable messages that are incorrectly intercepted by classic filters in a week). This threshold for unsolicited contact, along with his popularity as a target, is translated into the instantaneous value of his time. This, in turn, factors into the behavior of senders. Those with the most confidence in their ability to satisfy the recipient's needs do not fear the process. But senders without intimate knowledge of empowered recipients would never add risk to their message.

7.9 Does the system prevent 'zombies' from appropriating equipment? Who pays?

Yes. It is prevented by design and also by economic incentives...

Since 60-80% of spam is sent by zombies, critics object that an economic solution would hurt senders whose accounts get drained by infected PCs that send spam. Worse, novices and grandmothers would suffer the greatest harm because they don't have the skills of systems administrators to protect their machines. This is a very interesting critique, yet in 7.9.2, we demonstrate that it nicely illustrates the robustness of a pure economic solution.

7.9.1 Protection By Design

Funds at risk for individual senders are very small. Typically, $2 covers the risk of sending to other empowered users during a three day liability window, because the risk applies only to other risk-aware users and only those who have never sent mail back to the sender.

The mechanism adjusts the maximum daily liability to an amount equal to 300% of a rolling daily average. In the event that a sender issues more than 3x their normal daily mail volume to empowered users, they are asked to enter a 2nd level authentication.

7.9.2 Protection by Economic Incentive

Consumers and advertisers constitute a *two-sided network.* [6] In two-sided networks, one side gets a discount to bring them into the network. For example, Adobe offers a free reader to entice users to adopt their PDF format. Of course, they charge for the writer. Microsoft gives away Developer Toolkits to promote Windows applications but charges for Windows. Ebay gives buyers free access to auctions while charging senders.

In the case of spam Zombies, credit cards offer a good analogy. Merchants and consumers represent two sides of a network. If your card is stolen, your bank indemnifies you against fraudulent purchases over $50 as long as you report the problem. Why? Because the value of credit and debit card transactions in the US alone exceeds $1 trillion. If banks can skim a tiny fraction of this, they receive a lot of money. People won't use credit cards if it places them at risk, so it is profitable to insure them.

The same solution applies to Zombies. In exchange for a fraction of the spam bond proceeds, the ISP insures novices, and everyone else for fraudulent sending over $5, but only if they adopt an ISP's approved antiviral process. Even an institution with the skill to address the problem also has incentive to do it. The solution actually attacks spam at its roots and not just at the filtration level. Spammers will have a harder time getting into grandmother's machine because it's now actively protected. Meanwhile, advertisers gain a secure marketing medium and the ISP gains a fraction of the transactions. All parties to the network win, except the spammers.

7.10 Can it be effective before adoption?

The two-sided network effect applies here too. When an economically empowered sender and recipient meet, benefit accrues to everyone in the chain – except the spammer. Viral growth is a natural byproduct because positive results are easily observed by off-network users. In trials over two years involving tens of thousands of users, first-time recipients sense the intoxication of their passive power because of an unexpected reality: messages from strangers are suddenly very relevant.

7.11 What about desirable contact from those with insufficient resources?

Because of the way in which trust builds and risk expires, individual use can be covered for life by a one time $2 liability. This risk guarantee is even included in the cost of several anti spam products. They will only need to add cash if send thousands of messages to strangers in a short time or they irritate their regular contacts.

Commercial senders who prefer to avoid a large cash guarantee can build trust rapidly by spreading out their new "acquisition" contacts over several days. Alternatively, mailing list will guarantee the mailings of their clients while providing the added value of assured delivery, if they are certain that the message content is appropriate for their list.

7.12 Do ISP/ESPs have incentive to participate?

The acquisition of new customers is an ISP's largest expense. With a dramatic reduction in spam, ISPs saves money and generate satisfied customers. This translates into a low rate of customer turnover which, in turn, improves revenue and profit.

Because commercial senders quickly see economic benefit, some ISPs and ESPs will demand a micropayment from these senders in exchange for the risk-waiver test. We're already seeing this happen with AOL and Yahoo announcing their use GoodMail's CertifiedMail. It is an offer of preferential treatment in exchange for a cash payment.

But the ability to charge for something that has low overhead (testing for risk – a process that is free) is possible only if it ads competitive value. As competing ISPs and email services catch up, *it will be a competitive liability to lack the test*. Even for ISPs that continue to charge, the payment will be far lower than the risk warranty embedded within the message and it will apply only to first-time contact between strangers.

8. CONCLUSION

Will email no longer be free? It will remain free for you, your friends, your grandmother, and any marketer that respects you on your own terms. But the era of reaching out and irritating recipients is coming to an end. Every message that you receive will compensate you with pleasure, satisfaction or cash:

- The pleasure of receiving mail from a trusted friend
- The satisfaction of receiving mail with content that you find personally relevant and desirable
- An instantaneous cash transfer that actually rises to the true value of interrupting your otherwise pleasurable day

Recipients who take advantage of the payment/penalty mechanism quickly learn that cash warranties evaporate from their In Box (along with relevant contact from strangers – even if they are currently seeking information). That is, the economic feedback mechanism ensures that recipients who frequently "take cash from the table" are rarely offered cash.

In fact, with a pure economic mechanism, cash transfers are rare. The mechanism is not so much a *system of payment* as it is a *system of deterrence*. It effectively makes filters very smart by pushing the decision to filter all the way back to the sender.

He's not on my white list and the message content doesn't match my filters…

I get this back if the boss got no beef with the content of my message. Kapisch?

Footnotes

[1] Plan to Charge for Email Triggers Outcry. Feb 6, 2006. http://money.cnn.com/2006/02/06/technology/yahoo_aol_email/?cnn=yes

[2] AYRES, IAN & NALEBUFF, BARRY. *Want to Call Me? Pay Me!* Wall Street Journal, Oct 8, 2003

[3] In practice, the warranty collection rate is less than 1 in 500 and is almost always due to recipient error or the email being opened by someone other than the intended recipient. This is because senders who voluntarily risk a penalty generally have good information about their address list. They are supremely confident in the appropriate nature of their content to every recipient. Additionally, the mechanism punishes recipients who penalize relevant contact by substantially reducing the potential for delivery of relevant content in the future.

[4] LODER, THEDE & VAN ALSTYNE, MARSHALL & WASH, RICK. *An Economic Answer to Spam,* MIT / Boston U / U of Michigan, 2006

[5] WETZEL REBECCA. *Spam Fighting Business Models—Who Wins, Who Loses,* Business Communications Review, Apr 2004

[6] PARKER, GEOFFREY & MARSHALL VAN ALSTYNE Two-Sided Network Effects: *A Theory of Information Product Design,* Management Design Oct 2005

[7] www.vqme.com –Experiment with up to 5 email addresses without charge. Access clean email from any PC, even without software.