The Role of Public Policy in the Fight Against Spam

Jacob Scott
University of California, Berkeley
2004 WISE Intern
6 August 2004

Sponsored By

Table of Contents

Preface

About the Author

Jacob Scott, 21, was born in Oakland, California and will graduate from the University of California /Berkeley's School of Electrical Engineering and Computer Science in May 2005. He will spend the fall 2004 term as a research intern at Microsoft Research in Cambridge, England. Jacob's research interests focus on computer science theory and he plans to pursue a Ph.D. when he completes his undergraduate studies. His work experience includes an internship at Amazon.com and an undergraduate research fellowship at the National Institute of Standards and Technology. He has also worked as a research assistant to Professor Richard Karp at UC Berkeley and as a tutor in the university's department of computer science.

About WISE

Founded in 1980 through the collaborative efforts of several professional engineering societies, the Washington Internships for Students of Engineering has become one of the premier Washington internship programs. Its goal is to groom future leaders of the engineering profession who are aware of and can contribute to the important intersections of technology and public policy.

Acknowledgements

The author would like to thank IEEE-USA for providing the opportunity and resources that allowed him to pursue this project. Chris Brantley provided invaluable support and excellent advice, which contributed a great deal to the author's understanding of public policy. The guidance and knowledge of Professor Sharon Jones, the faculty

member in residence, played a large role in the success of the program, in general, and of the author's work, in particular. Finally, he would like to thank all those individuals, in both the public and private sectors, who were gracious enough to take time out of their busy schedules to discuss spam and related issues.

Executive Summary

In only a few years, spam has grown from a minor annoyance for a few people into a major problem for almost all those who use e-mail. This increase has been accompanied by increasing attention from the technology sector and the introduction of a wide variety of technologies promising to defeat it. Spam also receives growing exposure in the popular press as it becomes a daily issue for millions of Americans. More recently, spam has attracted the attention of the public sector, which, recognizing the varied and substantial costs associated with spam, has begun to explore ways to use public policy to combat it.

This paper examines how public policy can best contribute to the fight against spam. It begins with a general introduction to the spam ecosystem, highlighting the key factors that have allowed spam to become the problem it is today and its impact, both quantitatively and qualitatively. A more in-depth look at the technological context surrounding spam follows, which examines in more detail the weaknesses that have allowed spam to flourish, presents current anti-spam technologies and discusses their limitations. This background material on spam concludes with an overview of next-generation spam-fighting technology and how spammers are able to take advantage of technology.

Having covered the basics, the paper continues with a discussion of public policy. The CAN-SPAM Act, the first national anti-spam bill, is presented as a prime example of how the public sector, specifically, in particular, the US federal government,)approaches the fight against spam. CAN-SPAM is both presented and evaluated; emphasis is placed on the significant and controversial portions of the Act.

The tradeoffs inherent in drafting public policy and possible explanations for weaker parts of the bill are presented. The key conclusion here is that the legislation focuses on privacy and consumer rights, while it should, instead, address the more important issue of enforcement.

The paper concludes with recommendations on those public sector actions that would have the most significant impact on fighting spam. While there is an important role for the public sector in stewarding new anti-spam technologies and educating users on how they can reduce their exposure to spam, the paper concludes that increased enforcement of current laws (such as the CAN-SPAM Act, among others) against spammers will have the single largest impact. If such enforcement proves to be infeasible, the paper suggests that creating an individual right of action whereby those harmed by spam could bring legal action against spammers directly.

Introduction

Today, spam threatens the viability of e-mail as a communications medium. Just over a year ago, Federal Trade Commission (FTC) Commissioner Orson Swindle testified to Congress that "Spam is about to kill the 'killer app' of the Internet - specifically, consumer use of e-mail and e-commerce."[1] While this apocryphal vision has yet to be realized, the amount of spam spewed onto the Internet continues to increase. Anti-spam vendor Brightmail reports that 64 percent of e-mail it scanned in June 2004 – more than 100 billion messages[2] -- was found to be spam. In May 2004, America Online blocked up to 2.5 billion pieces of unwanted e-mail per day.[3] These messages are rarely legitimate marketing communications, given that e-mail users confront deceptive, fraudulent, and pornographic spam on a daily basis. An April 2003 FTC study of 1,000 pieces of spam found that two-thirds was, in some way, illegitimate and 18 percent was of an adult nature.[4]

Swindle's declaration that the nature and volume of spam entering consumer inboxes has begun to discourage e-mail's use is welldocumented. In a March 2004 Pew Internet survey, 29 percent of e-mail users reported reducing their overall use of e-mail because of spam. By comparison, in June 2003, 25 percent of users reported a reduced use of e-mail. [5] E-mail is widely viewed as the foundation of e-commerce and many other information technology-related innovations and if spam succeeds in undermining

---

[1] Orson Swindle. *Statement before the House Subcommittee on Commerce, et all*. Federal Trade Commission. June 11, 2003.
[2] Brightmail. http://brightmail.com/spamstats.html
[3] Ted Leonsis. *Testimony to the US Senate Committee on Commerce, Science & Technology*. America Online. May 20, 2004.
[4] *False Claims in Spam*. Federal Trade Commission. April 30, 2003. 2,10.
[5] Lee Rainie, Deborah Fallows. *The Impact of CAN-SPAM Legislation.* Pew Internet & American Life Project. March 2004. 1.

consumer confidence in e-mail, the economic and social impacts will be disastrous. In economic terms, market research firm The Radicati Group estimates that in 2003, spam cost businesses $20 billion worldwide.

Spam is at a tipping point. Within a very short time, it has been transformed from an annoyance into a grave problem. In the most basic analysis, spam represents the economic exploitation of an overly trusting technology, enabled by anonymity. There is no silver bullet that will solve the problem.

Technology certainly plays a central and crucial role in the fight to control spam. Without the filters that protect inboxes worldwide, e-mail would probably have been rendered unusable long ago. Nonetheless, the volume of spam continues to explode despite the wide array of technologies deployed against it. Bill Gates describes how spammers react to improved anti-spam tools. "Knowing that only a small percentage of their output will get past today's filters, spammers have responded by significantly cranking up the volume of emails they send," he says. "So networks are burdened with even more junk than before."[6]

The continuing burden of spam and the failure of technology alone to address the problem has brought this issue into the public policy arena. As both domestic legislation (including the CAN-SPAM Act of 2003) and recent international meetings of the Organisation for Economic Cooperation and Development (OECD) and the United Nations International Telecommunications Union suggest, legislation and enforcement, along with technological advances, will be required if spam is to be significantly

[6] Bill Gates. *Preserving and Enhancing the Benefits of Email – A Progress Report.* Microsoft. June 28, 2004.

curtailed. [7] However, the perceived failure of current anti-spam policies raises questions about whether and how the public sector can be effective in aiding technologists and industry in the anti-spam fight.

Given the nature of spam and the strengths and weaknesses of the technologies that both help and hinder it (and especially in light of recent industry efforts to require sender authentication to build accountability into the e-mail system), this paper proposes that policymakers' primary goals should be to pass legislation and, above all, establish vigorous enforcement of effective sanctions on spammers who continue to abuse the e-mail system for private gain. Enforcement is especially important at the international level, where issues of varying jurisdiction and legislative approaches create barriers to dealing with a problem that is global in scope.

Spam, Its Causes and Effects

Defining Spam

Developing an accurate and useful definition of spam is more complicated than it might appear. Every e-mail user recognizes spam when she or he sees it, but such notions may vary widely. Although the broadest definition would describe spam as a subset of unwanted e-mail, users' preferences vary with respect to the e-mail they would like to receive. As the Congressional Research Service notes, "the differences in defining spam add to the complexity of devising legislative or regulatory remedies for it."[8] Even

[7] The OECD had its first workshop on spam in February 2004, in Brussels France, with a second workshop planned in September 2004, in Busan, Korea. The ITU held a WSIS thematic meeting on countering spam in July 2004, in Geneva, Switzerland.
[8] Marcia S. Smith. *RL31953 - "Junk Email": An Overview of Issues Concerning Commercial Email and "Spam"*. Congressional Research Service. June 3, 2004. 4.

the Internet Research Task Force's Anti-Spam Research Group has not managed to decide on an all-encompassing definition of spam[9].

Despite confusion and disagreement on a precise definition, there is fairly widespread agreement that spam exhibits certain general characteristics. The OECD summarizes them well. First, spam is an electronic message. (For most purposes, this may be restricted to e-mail, but other methods of delivering spam do exist, including the Short Messaging Service, or SMS).

Second, spam is unsolicited. If the recipient has agreed to accept a message, it is not spam. However, how and when such consent is given may not be clear, especially when a pre-existing relationship exists between the sender and recipient.

Third, spam is sent in bulk. This implies that the sender distributes a large number of essentially identical messages and that recipients are chosen indiscriminately.

These three traits define Unsolicited Bulk E-mail (UBE). If a fourth is added -- that spam must be of a commercial nature -- the resulting class of messages is referred to as Unsolicited Commercial E-mail (UCE).[10]

Even defining spam as UCE instead of UBE is a controversial choice. Many consumers view unsolicited bulk mailings of a religious or political nature, for example, as spam. In June 2004, Wired.com described a 2004 incident in which "E-mail users around the world got a rude awakening Thursday when a spammer flooded their inboxes with nationalist, borderline-racist propaganda in German."[11] Some could indeed argue that these e-mails met the definition of spam, although they were not UCE. Furthermore,

[9] *FAQ*. Anti-Spam Research Group. May 21, 2004.
[10] *Background Paper for the OECD Workshop on Spam.* OECD. January 22, 2004. 7-8.
[11] Amit Asaravala.*German Spam Floods Inboxes.* Wired News. June 11, 2004.

it may be difficult for automated filters to determine whether an e-mail has commercial content, so in this context, spam may, nonetheless, end up being defined as UBE.

However, advantages do exist to limiting spam to UCE. First, there may be legal complications, such as First Amendment concerns, with legislation that tries to regulate UBE[12]. Second, since in this case, spam is , by definition, commercial speech, it can be addressed on the economic front by limiting its profitability.

While acknowledging these current controversies, spam will be defined here to include UCE and certain types of UBE, like those containing viruses, worms, or Trojan Horses. This definition closely mirrors those found in worldwide anti-spam legislation, while also encompassing a large segment of UBE.

There is a significant qualitative difference between the spam sent by so-called outlaw spammers and that sent by more legitimate e-mail marketers. Generally, 'spam' and 'spammer' in this paper refer to the worst spam (that is, deceptive, pornographic, and fraudulent) sent by the worst spammers because they represent the crux of the spam problem. While it would be unfair and inaccurate to attribute the same behaviors to all senders of UCE, a reputable sender or a Fortune 500 company is not automatically exempted from being a spammer.

The Problem

Spam has negative impacts for consumers, businesses, Internet Service Providers (ISPs), legitimate e-mail marketers and virtually anyone else who uses e-mail for any reason. However, listing the burdens spam imposes does not provide a full picture of the

[12] *Additional EFF Comments for CAN-SPAM Act ANPR.* The Electronic Frontier Foundation. April 9, 2004.

problem it creates. Rather, the reasons spam is sent help explain why it has become such a severe problem. The next section addresses the causes and effects of spam.

Causes of Spam

Spam has become a serious problem because it is profitable and can be almost completely anonymous. In this sense, it is a problem of the commons. Spammers are polluting a common resource (e-mail) for their own short-term profit and are protected from being held accountable for their behavior because of the difficulty of identifying them.

The profits associated with spamming are based on the same principles that make so-called legitimate e-mail marketing profitable; that is, because some people respond to spam by buying the products advertised. According to Pew Internet, five percent of e-mail users have ordered a product or service offered in an unsolicited e-mail.[13] More generally, the Direct Marketing Association (DMA) estimates that an excess of $19 billion was spent in response to commercial e-mails in 2003.[14] This represents a large and growing market for spammers.

In addition, the low cost of e-mail as a communications medium virtually guarantees profits. The marginal cost of sending an e-mail message has been estimated at .05 cents[15], which is essentially negligible. Thus, spam can be profitable even at an extremely low response rate. As a representative of the Electronic Information Privacy Center (EPIC) testified to Congress, "the marginal cost of each additional electronic message is essentially zero. Therefore, spammers are as likely to send to a million users

[13] Lee Rainie, Deborah Fallows. *The Impact of CAN-SPAM Legislation.* Pew Internet & American Life Project. March 2004. 5.
[14] Peter A. Johnson. *Preserving The Promise of The E-mail Marketplace.* Direct Marketing Association. March 31, 2004. 11-12.
[15] Paul Judge CipherTrust.

as they are to [send to] a thousand."[16] FTC Chairman Timothy Muris has referred to the case of "a bulk emailer who testified that he could profit even if his response rate was less than 0.0001 percent."[17]

Because of technological obstacles associated with e-mail's infrastructure [details are discussed in a later section], it is notoriously difficult and time-consuming to trace a piece of spam back to the individual or group responsible for its transmission. Spammers are hard at work making tracing even more difficult, including by moving their mail through offshore ISPs. Steve Linford, director of Spamhaus (a major blacklist operator), reported in June 2004 that up to 70 percent of spam was being sent from Chinese ISPs by American spammers.[18] Such actions complicate enforcement and allow spammers to remain anonymous, thereby allowing them to continue spamming with impunity.

In summary, the root causes of spam are consumers' willingness to purchase products advertised in unsolicited e-mail and the inexpensive and anonymous communications mechanism that e-mail provides. Its effects, however, are extremely wide reaching.

Effects of Spam

Although it is difficult to quantify the inconvenience and burden of spam on those who receive it, several recent studies do attempt to assess that cost. Ferris Research, an e-mail focused market research firm, estimates that in 2003, spam cost US-based organizations $10 billion, including computing resources, administrative and helpdesk

[16] Marc Rotenberg. *Testimony and Statement for the Record.* Electronic Privacy Information Center. May 21, 2003. 2.
[17] . Timothy Muris. *Remarks, Aspen Summit*. Federal Trade Commission. August 19, 2003.
[18] Graeme Wearden. *Russia and China 'behind current spam deluge'.* ZDNet UK. June 9, 2004.

personnel time and lost productivity[19]. A European Union report estimates annual worldwide costs to individuals alone at over 10 billion euros[20].

For the consumer, the spam deluge creates a significant problem -- that of the *false positive*. These are the pieces of genuine e-mail that are falsely identified as spam, whether by a filter or by accidental deletion. Research conducted by anti-spam vendor Goodmail Systems suggests that 67 percent of e-mail users have lost e-mail because of this false positive problem[21]. As more and more forms of communication are delivered by email, losing such messages can have serious repercussions if, for example, they involve a bill or a legal notice. In comments to the FTC, Virginia Assistant Attorney General Prosecutor Russell McGuire referred to the latter situation, noting that, "you might be getting notice of a particular docket hearing… Saying, 'My filter blocked it,' is not going to be an excuse when you get into court."[22]

Consumers also confront deceptive, fraudulent and pornographic spam. A Pew Internet study describes the impact of adult-oriented spam in a section titled, "There is a special place in Hell for pornographic spam." In the report's conclusion, the author, Deborah Fallows, notes, "In nearly every measure we tested, pornography soared to the top as the most offensive, objectionable, destructive type of spam."[23]

Another consumer concern is the growing practice of *phishing*. Phishing refers to a type of spam that uses e-mails and websites designed to resemble those of legitimate organizations (such as banks, ISPs, or government agency) to trick recipients into

---

[19] Ferris Research. http://www.ferris.com/url/spam.html.
[20] Serge Gauthronet, Etienne Drouard. *Unsolicited Commercial Communications and Data Protection*. Commision of the European Communities. January 2001. 67.
[21] *FAQ*. Goodmail Systems. 2004.
[22] *Do Not E-Mail Registry Meeting*. Federal Trade Commission. March 10, 2004. 46.
[23] Deborah Fallows. *Spam: How It Is Hurting Email and Degrading Life on the Internet.* Pew Internet & American Life Project. October 22, 2003. 29, 41.

disclosing sensitive information like Social Security numbers, credit card numbers and bank account information.[24] The Anti-Phishing Working Group reports that well over 250 unique phishing attacks were recorded every week in May 2004.[25] Also in May, the Gartner research firm reported that banks and credit card companies had suffered $1.2 billion in direct loses to phishing during the previous year.[26]

Businesses face reduced worker productivity and increased information technology costs as users delete spam from their mailboxes and system administrators struggle to limit its effects on corporate servers. Market research firm Nucleus Research estimates the costs of lost productivity alone to be over $1,900 per person per year, more than double the amount reported in 2003[27]. A thriving sector for business oriented anti-spam solutions has arisen to address these concerns, but while they may have a good return-on-investment (the cost of using such products is less than that of subjecting employees to unfiltered spam), they are not free.

Spam places a heavy burden on ISPs, the backbone of the information superhighway. Bill Gates has said that Hotmail blocks an average of almost 3 billion pieces of spam every day.[28] Spam affects ISPs with respect to bandwidth used in its transmission, disk space used to store it, staff resources (both legal and technical) to maintain the effectiveness of anti-spam measures and pursue enforcement, and customer service resources devoted to educating users on spam's dangers and responding to

[24] http://www.ftc.gov/bcp/conline/pubs/alerts/phishingalrt.htm
[25] *Phishing Attack Trends Report*. Anti-Phishing Working Group. June 22, 2004. 1.
[26] Avivah Litan. *Phishing Victims Likely Will Suffer Identity Theft Fraud*. Gartner. May 14, 2004.
[27] *Spam: The Serial ROI Killer (Research Note E50).* Nucleus Research. June, 2004. 1.
[28] Bill Gates. *Preserving and Enhancing the Benefits of Email – A Progress Report.* Microsoft. June 28, 2004.

complaints and questions. Finally, ISPs must remain vigilant to ensure that their own subscribers are not engaged in spamming.

Spam presents legitimate online marketers with yet another set of problems. The most important is that the very large volume of outlaw spam may drown out their communications. This is one reason why many marketing trade groups strongly supported the CAN-SPAM Act, which industry representatives believed would distinguish legitimate commercial e-mail from unlawful spam[29].

Another issue is a type of malicious attack known as a "Joe job"[30]. This refers to a situation in which a spammer falsifies various parts of an e-mail to forge the sender's identity. The apparent sender may have had nothing to do with the spam message. While "Joe jobs" are sometimes defined as requiring specific intent to harm the falsified sender, a large amount of spam has a forged 'From' line, to similar effect. These incidents may harm the reputation of the misidentified sender, inundate that sender's mailboxes with 'bounce' messages or generate misdirected complaints from the spam's recipients.

As this analysis demonstrates, spam poses a wide range of problems. It is a productivity problem, costing businesses and ISPs billions of dollars. It is a fraud and identity theft problem to the tune of billions of dollars. Finally, it threatens to stifle information technology on a larger scale, from parents worried about their children receiving pornographic spam to consumers losing trust in and abandoning e-mail altogether. It is not, however, an unrecognized problem. The technology industry has recognized the demand for solutions to spam and has responded vigorously by

[29] Press Release, Direct Marketing Association. December 16, 2003.
[30] The name for this term originates from its first appearance, launched against joes.com. See http://www.joes.com/spammed.html

committing resources to research and development. A host of excellent anti-spam measures have been implemented.

The Role of Technology in Addressing Spam

The State of E-mail

Imagine the following changes to US Postal Service operations. First, postage, envelopes and paper are free. Second, the postmark has been abolished, so when addressees receive a piece of mail, there is no way to know whether it was sent from Maine or Hawaii because senders cannot be forced to put a valid return address on their outgoing mail.

On the one hand, this postal service would be very convenient for most users because it would be free and it would be easy to send letters to anyone in the world. Unfortunately, such a free and anonymous system would also quickly result in mailboxes clogged with unimportant content, sent by those taking advantage of a free and anonymous system. This is close to how e-mail operates today.

The design of the Simple Mail Transfer Protocol (SMTP), which defines how e-mail moves through the Internet, does little to prevent spammers from exploiting the system. The specification that defines SMTP (a Request For Comments, or RFC) has not changed since it was originally formalized in 1982[31] and provides no mechanisms by which to verify the headers or routing information that make up e-mail metadata, which includes information about the sender's identity and the route the message traveled across the Internet.

[31] RFC 821 was superceded by RFC 2821 in 2001, but no functionality was added or changed

Without additional mechanisms (for example, cryptographic digital signatures), it is impossible to verify whether the purported e-mail sender is the actual sender or whether the email arrived by the path it claims to have taken. Currently, the only piece of information that an e-mail recipient can trust to be accurate is the sending server's Internet Protocol (IP) address.[32] These shortcomings have led scientists and engineers to suggest that SMTP be abandoned in favor of a newer protocol (recreating e-mail from scratch), despite the cost.[33]

Spam Filtering

SMTP, along with a handful of other closely related protocols[34], provides the framework for spam-fighting technologies. To date, the focus of such technologies has been on filters. Generally, filters do not modify SMTP but instead operate on e-mail after it has been received (either at the server or at the user's mailbox). Based on content, IP address, behavior or any other information they can glean about the e-mail, filters decide whether a particular message is likely to be spam. They either reject it immediately or route it away from the inbox and into a special spam folder. While the primary goal of filtering is to lessen the impact of spam on e-mail users, some hope that if enough spam is filtered, it will reduce spamming's profitability.

The techniques used to filter e-mail are extensive and diverse. They range from monitoring the real-time behavior of e-mail servers to collaborative methods like Vipul's Razor. In July 2004, anti-spam vendor Postini reported that it was blocking 53 percent of

[32] *Technology and Policy Proposal*. Anti-Spam Technical Alliance. June 22, 2004. 17.
[33] Paul Festa. *End of the road for SMTP?* C|Net News.com. August 1, 2003.
[34] RFC 2822 defines the format of e-mail messages, while RFC 2476 defines an e-mail submission protocol that can be used in concert with SMTP

all e-mail "based on the real time behavior of the spammers' IP addresses."[35] Vipul's Razor relies on contributions from its users to establish a distributed and constantly updated catalogue of spam in propagation that email clients consult to filter out known spam.[36] Customized filters allow users to utilize their own definition of spam to determine what to block.

Filters aim for high effectiveness and high accuracy. High effectiveness refers to the volume of spam blocked, while high accuracy means that the filter's error rate for misidentifying real e-mail as spam, thus generating false positives, is low. A number of robust commercial and open source products have been developed in recent years. Brightmail, for example, advertises that its enterprise-oriented Anti-Spam 6.0 product blocks 95 percent of spam with a 99.9999 percent accuracy rate, generating less than 1 false positive for every million spam messages blocked.[37] In March 2004, Bill Yerazunis, the author of CRM114, an open-source spam classifier, reported 2004 that his filter had made only one mistake in over 8,700 message classifications, making it ten times more effective at detecting e-mails than he was as an "unassisted human."[38]

Beyond accuracy and effectiveness, good spam filters will try both to reduce the amount of required user interaction (for the greatest possible transparency) and the bandwidth, storage, and computational overhead that the high volume of spam imposes on e-mail users and service providers.

[35] Press release, Postini. July 7, 2004.
[36] http://razor.sourceforge.net/
[37] http://www.brightmail.com/enterprise-as-benefits.html
[38] http://crm114.sourceforge.net/

Combination is also powerful. Many spam-fighting technologies work well in concert.[39] Such applications can reinforce the strengths and reduce the weaknesses of individual approaches. Individual technologies are also often flexible in their point of application.

The Limits of Filtering

Filtering as a general technique has two major vulnerabilities. The first is that of volume, which was mentioned in the introduction. Even a highly effective filter can be rendered useless if it is just as easy to send millions of e-mails as it is to send thousands. A background paper for the ITU spam summit describes the problem. "[O]ften described as an 'arms race' between filter makers and spammers, today's situation can also be seen as a chicken and egg problem: spammers send more messages because there are more filters, and there are more filters because spammers continue to send more messages."[40] Although clearly preferable to spam that makes its way to the inbox, blocked spam messages do present bandwidth and computational costs.

The second problem for filters is that spam is not a static problem but a dynamic one. In addition to spewing out more e-mail, spammers continue to adapt their methods to evade current spam-blocking technology, which in turn forces the anti-spam community to update its tools.[41] E-mail users are far more tolerant of false negatives (spam in their mailbox) than false positives (real e-mail discarded). The more closely a piece of spam resembles a legitimate e-mail, the less effective filtering becomes. In this regard, the most effective methods to circumvent filters rely on deceit and fraud. That is,

---

[39] SpamAssassin, a popular open-source e-mail filter, lists five separate technical approaches to spam identification. http://spamassassin.apache.org

[40] Unspam WSIS

[41] Ingrid Marson. *Brightmail's new software tackles zombies*. ZDNet UK. July 1, 2004.

such technologies are most easily thwarted by the most serious spammers, those who do not hesitate to hide their identities or send deceptive messages.

Stopping Spam at the Source

In addition to filtering, significant resources have been devoted to tools to prevent spammers from using ISPs' resources. These technologies represent an effort on the part of ISPs to ensure that their e-mail servers are not abused. In effect, they are telling spammers, "Not in my neighborhood." Most are policy- and security-based responses, rather than cutting-edge technologies.

The degree to which spam can be reduced by simple ISP policy changes is remarkable. In June 2004, broadband ISP Comcast found that spam originating from its networks dropped 35 percent in less than a month after it blocked outgoing e-mail traffic that did not originate from its mail servers. This action prevented its customers' connections from being hijacked by spammers.[42] During that same month, the Anti-Spam Technical Alliance, an industry group whose founding members include America Online and Microsoft, released a best practices document that included ten recommendations that ISPs could use to help reduce spam.

ISP action to reduce spam is generally well received, but there are tradeoffs. For example, ASTA suggests rate-limiting outgoing e-mail[43]; that is, limiting the number of e-mails that consumers are allowed to send in a fixed amount of time. This action could impact legitimate users who may run high-traffic mailing lists on their consumer accounts. Further, although actions to secure networks and servers are important, the Internet's large size means that so far, spammers have always been able to look for and

[42] Jim Hu. *Comcast reports 35 percent decline in spam*. C|net News.com. June 29, 2004.
[43] *Technology and Policy Proposal*. Anti-Spam Technical Alliance. June 22, 2004. 12.

exploit a weaker link in the chain. Because spammers can conceivably run their own ISPs and because those ISPs that provide services to spammers also profit, that weakest link is unlikely to vanish soon.

Addressing the Root Causes of Spam with Technology

Anti-spam technology addressing the root causes of spam is only now moving from the research stage to deployment and follows two major approaches. The first is sender authentication, which allows e-mail recipients to confirm that senders of e-mail messages are not forging their identity.[44] In addition to providing recipients with another way to judge the legitimacy of e-mail, sender authentication protects e-mail senders from Joe-jobs and phishing attacks.

Reflecting their maturity, the two major projects in this area, DomainKeys and Sender ID, have both recently submitted Internet Drafts to the Internet Engineering Task Force (IETF), the Internet's main standards body. Meng Weng Wong, an author of Sender ID, discussed deployment in a June interview about Sender Policy Framework (SPF), a central component of Sender ID. "We expect adoption to pick up exponentially; according to some estimates, the number of sites checking SPF doubles every three weeks," Wong said.[45]

The second major push is for cost-shifting technologies that raise the spammers' marginal cost of sending e-mail without destroying the usefulness of e-mail for everyone else. Microsoft's Penny Black project proposes that unsolicited e-mail carry a *computational* payment. The approach is based on the notion that if the recipient does not

[44] Current approaches generally authenticate the domain, rather than the full email address, of the sender. This provides some privacy while ensuring that if forgery does occur, both the actual and alleged senders will be authorized senders for the domain in question.
[45] *An Interview with the Lead Developer of SPF – Part I*. Circle ID. June 29, 2004.

already know a sender who wishes to send the recipient e-mail, the sender must prove that it has expended a certain amount of effort to reach that particular recipient with a particular message. [46] The ideal cost of such a computational payment (which would be tendered, for example, by solving a puzzle of established difficulty) would not inconvenience normal users of e-mail, but would significantly raise marginal costs for spammers, forcing them to purchase more servers if they wanted to continue sending millions of messages a day.

There are also proposals to introduce monetary payments to deter spammers, either by "e-mail stamps" or by various sorts of bonds. The most well known is probably IronPort's Bonded Sender program. Bonded Sender allows commercial e-mailers to post a bond as proof of their good intentions. In exchange, they receive some level of assurance that their mail will be delivered to recipients rather than rejected by ISP spam filters. MSN's Hotmail service worked closely with IronPort on Bonded Sender's development, and the two have reported some initial success.[47]

Another proposal is the Attention Bond Mechanism (ABM), which allows e-mail users to decide how much the inconvenience of reading a piece of spam is worth. It then requires those sending them e-mail to post a bond in that amount that would be cashed if the recipient concluded that the message was spam. Proponents claim that allowing individuals to determine the amount of the bond makes ABM a more superior mechanism than even a perfect filter because it reimburses e-mail recipients for their wasted time.[48]

[46] http://research.microsoft.com/research/sv/PennyBlack/
[47] Press release, IronPort. May 5, 2004.
[48] http://www.eecs.umich.edu/~tloder/one_pager.html

Drawbacks

A major obstacle facing technologies that would change how e-mail operates is that significant resources are invested in the current e-mail status quo. Despite its failings, SMTP is deeply embedded in millions of servers, companies, and ISPs worldwide. The simpler and less costly the anti-spam solution, the more widely it will be adopted.

The technologies discussed in this section rely, in large part, on wide adoption. Meng Weng Wong refers to this as the fax effect (the more people who have faxes, the more valuable a fax machine is to someone who would buy one). He notes that, "for the fax effect to fully manifest, a significant majority of legitimate e-mail on the Internet should be covered by Sender ID."[49] Further, these changes imply e-mail's transformation from an open, free, and anonymous communications medium into something else. Because these characteristics are widely recognized as the basis for e-mail's tremendous success, there is considerable concern about how far these changes will go. Members of the Internet community have been particularly opposed to, and have vigorously resisted, introducing monetary payments into the general e-mail ecosystem.

Sender authentication and cost-shifting systems also have specific weaknesses. Most sender authentication techniques rely on the Domain Name System (DNS) to store authentication records (the IETF has a working group on this issue[50]). Identity and authentication is then tied to the sender's domain, which means that if the system is to be effective, domain registries must have accurate information.

Payments systems, whether computational or monetary, face even more hurdles. First, many of these systems require whitelisting (that is, exempting legitimate

[49] Meng Weng Wong. *Behind The Curtain: An Apology for Sender ID*. Pobox.com. June 23 2004. 8.
[50] http://www.ietf.org/html.charters/marid-charter.html

individuals from the payment requirement) to reduce the number of total payments made. Secure whitelisting depends, in large part, on sender authentication. Another problem that payment systems face is mailing lists. One-to-many communication requires a substantially larger payment, assuming some whitelist scheme is not in place. Monetary systems rely on huge infrastructure upgrades and on overcoming the stigma associated with introducing money into the e-mail equation. Computational payments have been challenged in light of the large numbers of "spam zombies" (hacked consumer computers turned into spam servers) to which spammers have access and which they could use to make such payments. [51]

Using Technology to Spam

Spammers also take advantage of cutting edge technology at all levels of the spamming process. Technology aids spammers in three key ways. First, it finds valid e-mail addresses to which to send spam. Second, it helps spam avoid being recognized as an unwanted message and, thus, escape filtering. The third involves methods that provide spammers a means to hide their tracks and remain anonymous.

Address generation provides a particularly useful introduction to how spammers use technology. Spammers rely on two techniques to obtain most of their e-mail addresses: web harvesting and directory attacks (sometimes called directory harvesting attacks). Web harvesting software crawls the Internet just as all search engines do. However, it also scours the pages it views for e-mail addresses to which spammers later send mail. In the summer of 2002, the Center for Democracy and Technology launched a six-month project to discover where spammers were finding e-mail addresses. They

[51] Munir Kotadia. *Microsoft's anti-spam plan 'hijacked by zombies'*. ZDNet UK. June 4, 2004.

concluded that, "By an overwhelming margin, the greatest amount of spam we received was to addresses posted on the public Web."[52]

Spammers launch directory attacks by using automated software to connect to an e-mail server and then trying to send e-mail to thousands or millions of automatically-generated addresses, usually in some sort of alphabetical or random order. The addresses that the server does not reject will almost certainly be valid ones that the spammer can record and send e-mail to at a later date. Directory attacks have grown significantly, with Postini reporting 15 million directory attacks against its customers in April 2004.[53]

Of course, spammers do not always need to rely on technologically sophisticated techniques to find addresses of potential e-mail recipients. Observers believe that a thriving wholesale market exists for valid e-mail addresses. In June 2004, for example, two men (including an AOL software engineer) were arrested for stealing and selling a list of 92 million AOL members' screen names (e-mail addresses). The entire list was sold to other spammers for up to $52,000.[54]

The Failures and Successes of CAN-SPAM

In 1997, the state of Nevada passed the nation's first law against spam. Over the next six years, spam exploded and other states, including Colorado, Oregon, and Washington, began to pass similar legislation. Anti-spam bills began to be introduced in Congress in 1999. However, the most comprehensive legislative effort to date to address these issues emerged from the Senate in 2003.[55]

---

[52] http://www.cdt.org/speech/spam/030319spamreport.shtml
[53] Press release, Postini. May 4, 2004.
[54] Erin McClam. *Two arrested in AOL spam scheme, prosecutors say*. Associated Press. June 23, 2004.
[55] http://www.spamlaws.com

On December 16, 2003, President Bush signed into law S. 877, the *Controlling the Assault of Non-Solicited Pornography and Marketing Act of 2003*, or CAN-SPAM Act. The CAN-SPAM bill passed the Senate on October 22 by a 97-0 vote and the House on November 21 by a 392-5. As the first federal law to directly address spam[56], the Act became effective on January 1, 2004. Six months later, anti-spam vendor MX Logic reported that the average compliance rate for commercial e-mail during that period was 2.3 percent, falling from 3 percent in January 2004 to 1 percent in June 2004.[57] This led the law's critics to ask whether in this context, "CAN SPAM" meant the ability to send spam rather than the command to dispose of it. Proponents argue that it is too early to evaluate the law's success given the time required to implement enforcement and assess compliance. They also cite Congress' decision to give the FTC two years to draft a report on the Act's effectiveness to support their argument.[58]

The motivation for the CAN-SPAM Act is presented out in the Congressional Determination of Public Policy, in which Congress concluded that "senders of commercial electronic email should not mislead recipients as to the source or content of such mail; and recipients of commercial electronic mail have a right to decline to receive additional commercial electronic mail from the same source."[59] This represents a forceful approach to dealing with forgery and deception in e-mail sent by outlaw spammers, but a much less rigorous tact to dealing with privacy.

[56] "The CAN-SPAM Act defines "spam" as certain types of commercial e-mail. Thus the use of the term in this section should be seen as identical to that found in Section 3(2) of the Act. The differences between the CAN-SPAM Act definition and those discussed in *Defining Spam* are not large.
[57] Press release, MX Logic. July 7, 2004.
[58] *CAN-SPAM Act of 2003*. Section 10.
[59] *CAN-SPAM Act of 2003*. Section 2(b).

The Opt-in Opt-out Debate

Viewed one way, spam laws are primarily an issue of consumer protection and electronic privacy. The key question in such a context is whether marketers must receive affirmative consent from consumers before making contact with them (opt-in) or whether they are allowed to make "preemptive" contact with consumers if they provide a method for consumers to reject further contact (opt-out). The opt-out method is the weaker and that is the one the CAN-SPAM Act adopted.

Privacy and consumer advocates are classically on one side of this debate, with the e-mail marketing industry on the other. In May 2004 testimony before the Senate Commerce Committee, a Consumers Union representative offered an argument against opt-out. "Imagine that you put a 'do not solicit' sign at the front door of your home, and every company in the world could only ring your doorbell once, at which point you could tell the salesperson not to bother you anymore…," the representative noted. "This is an absurd burden to place on people."[60] Spam is clearly less intrusive than a solicitor or a telephone call, but as the importance of electronic communication (such as instant messaging and Voice over IP) continues to grow, opt-out sets a bad precedent.

Perhaps the most significant problem with opt-out is a practical one. It is widely believed that taking advantage of opt-out measures allows spammers to confirm a valid e-mail address and may even generate more spam. While a 2002 FTC study found that opting out did not do so, [61] Pew Internet reports that "The 'remove me' function is now confusing and untrustworthy."[62] While such a problem may be moot when only a

[60] James Guest. *Testimony of James Guest*. Consumers Union. May 20, 2004. 3.
[61] Timothy Murris. *Testimony to US Senate Committee on Commerce, Science, and Transportation*. Federal Trade Commission. May 20, 2004.
[62] Pew Internet 42.

negligible amount of e-mail complies with CAN-SPAM in the first place, it speaks to the need for consumer trust in e-mail and its senders before such measures will help reduce spam.

E-mail marketers oppose opt-in because of the economic losses associated with curtailing opt-out marketing. They argue that such provisions are unenforceable and serve only to punish the legitimate marketers who would comply, while outlaw spammers and other bad actors would continue to spam. Available statistics on CAN-SPAM compliance support the credibility of such theories.

The reality of spam is that today, stopping deception and fraud outweighs privacy considerations. Judging by volume or content, the worst spam is sent by outlaw spammers whose practices and methods flaunt the CAN-SPAM Act regulations. So while current anti-spam laws provide less privacy protection for e-mail users than consumers receive, for example, in their telephone communications (in the form of the very successful Do Not Call Registry), and while advertising that complies with the CAN-SPAM Act can often still be considered spam, the most problematic spam is so bad that it renders such distinctions moot.

Do Not Email Registry

A compromise between opt-in and opt-out in the form of a Do Not Email Registry (similar to the Do Not Call Registry) was included in the CAN-SPAM Act, which requires the FTC to issue a report on the viability of such a system. Ideally, a Do Not Email Registry would give individual consumers the right to choose opt-in without subjecting marketers to a global opt-in system. Thus, marketers would have to ensure that they did not send to any e-mail addresses on the Registry.

In June 2004, the FTC issued its report, concluding that "a National Do Not Email Registry, without a system in place to authenticate the origin of email messages, would fail to reduce the burden of spam and may even increase the amount of spam received by consumers."[63]

The FTC cites two main problems with the Do Not Email Registry approach. First, like many of the provisions in the CAN-SPAM Act, violations of the Registry provision would be extremely difficult to enforce in today's e-mail environment. Second, as noted by the Email Service Provider Coalition (ESPC) in testimony to the FTC, "If I have a list and I want to send a mail, and you want to tell me not to mail certain people on it, you have to tell me who not to mail it to."[64] The apparently unavoidable distribution of e-mail addresses to marketers makes it highly likely that they would eventually leak to those who would abuse them. An additional concern suggests that the success of the Do Not Call Registry and the likely failure of a Do Not Email Registry would create unreasonable consumer expectations and generate even more distrust in e-mail.

Short of opt-in, which is widely considered politically infeasible, a Do Not Email Registry is likely the only additional protection to be offered to consumers.The FTC has made a stronger enforcement environment a prerequisite for reconsidering the Registry. This further demonstrates the focus on enforcement over privacy and consumer rights.

Disallowing Abusive Practices

The core of the CAN-SPAM Act is language that clearly prohibits the deceptive, misleading, and fraudulent behaviors in which legitimate senders of e-mail would never engage in but that are often exhibited by outlaw spammers. Additions to Title 18 of the

[63] *National Do Not Email Registry: A Report to Congress*. Federal Trade Commission. June, 2004. i.
[64] *National Do Not Email Registry: A Report to Congress*. Federal Trade Commission. June, 2004. 22.

US Code make it illegal to use hacked zombies to send spam or fraudulently register for IP addresses, e-mail accounts, or domain names used to send spam.[65] It is also unlawful to materially falsify header information (including both routing information and sender e-mail address) or to use a deceptive subject when sending UCE.[66] The practice of using web-harvesting or dictionary attacks to find e-mail addresses, as discussed earlier, are considered aggravated violations.[67]

Requirements for Senders

The CAN-SPAM Act imposes three requirements on the content of a commercial e-mail: clear notice that the message is an advertisement; the opportunity to opt-out and a mechanism by which to do so; and, the sender's valid physical postal address.[68] However, it is not necessary that these "clear and conspicuous" pieces of information be in any machine-readable form. That is, there is no short cut provided that would allow filtering or opt-out responses to compliant messages to be sent automatically by a software program instead of by a human being. Indeed, the FTC is expressly forbidden from issuing rulemaking that would "establish a requirement… to include specific words, characters, marks, or labels"[69] in UCE, although it is charged with issuing a report on the issue to Congress within 18 months.[70] Interestingly, the FTC was allowed to, and in April 2004 did issue, a rule requiring that a subject line prefix --"SEXUALLY-EXPLICIT:"– be required at the start of the subject line of all sexually oriented commercial e-mail.[71]

---

[65] CAN-SPAM Act of 2003. Section 4 (a)(1) (Chapter 47, Title 18 of USC, 1037 (a))
[66] *CAN-SPAM Act of 2003*. Section 5 (a)(1,2).
[67] *CAN-SPAM Act of 2003*. Section 5 (b)(1).
[68] *CAN-SPAM Act of 2003*. Section 5(a)(5).
[69] *CAN-SPAM Act of 2003*. Section 13(b).
[70] *CAN-SPAM Act of 2003*. Section 11(2).
[71] *16 CFR Part 316: Label for E-mail Messages Containing Sexually Oriented Material; Final Rule*. Federal Trade Commission. April 19, 2004.

The argument against such subject line filtering in general (for example, an "ADV:" prefix) is enforcement-related. Under that scenario, bad actors and outlaw spammers (who are believed responsible for the vast majority of spam, by volume) will flout the law, legitimate marketers will be put at a disadvantage and consumer expectations of a simple solution to spam will be dashed. Nearly everyone, including marketers, agrees that adult-oriented e-mail constitutes a special case. No one who objects to it should be subject to seeing it, especially because e-mail addresses contain no age-identifying information. It is not surprising that adult-oriented spam shows higher CAN-SPAM compliance rates than the average spam message. MX Logic found that shortly after the FTC's rulemaking went into effect in May 2004, compliance was 15.3%.[72]

Closing the Loopholes

The CAN-SPAN Act recognizes how difficult it is to find and prosecute spammers by retracing their Internet footsteps, partly because spammers can move offshore so easily. To address these problems, the Act makes it illegal for anyone to "promote, or allow the promotion of, that person's trade or business"[73] via spam,[74] if they had some knowledge of the spam, profited, or expected to profit, from the spam, and took no action to stop, or attempt to stop, the spam. (The drafters hoped this wording would prevent "Joe jobs.") This allows the FTC [75] to "follow the money" and pursue those profiting from spamming. It is often much easier to track financial records like credit cards than e-mail communications because long-established procedures exist for

[72] Press release, MX Logic. June 9, 2004.
[73] *CAN-SPAM Act of 2003*. Section 6(a).
[74] Technically, commercial e-mail violating Section 5 (a)(1)
[75] Only the FTC is allowed to enforce Section 6. See Section 6(c).

determining fraud. Also, this language clearly states that spamming is not a reputable business practice and prevents disreputable advertisers from circumventing the Act by hiring others to spam and then claiming that they were not directly involved in the subsequent violations.

Penalties

The civil and criminal penalties levied against violators of the CAN-SPAM Act are considered to be very strong. The sections of the Act added to Title 18 and the aggravated violations are considered felonies and carry up to five-year sentences, and are prosecuted by the Department of Justice (DoJ). The FTC enforces the CAN-SPAM Act as if violations were unfair acts or practices under the FTC Act,[76] with each violation subject to a fine of up to $11,000 each.[77] There are no caps placed on the total amount of the fines the FTC can seek. Both State Attorneys General and ISPs can file civil suits under the CAN-SPAM Act and the statutory damages for violations concerning materially false and misleading header information are not subject to any cap.[78] In short, the Act makes it possible to bankrupt spammers. They also risk significant prison time for the worst violations.

Enforcement

Anti-spam enforcement is the key to preventing outlaw spammers from abusing e-mail. Those who have no respect for the law and are most intent intent on reaching consumers are responsible for the most offensive and problematic spam. They are the

[76] *CAN-SPAM Act of 2003*. Section 7(a).
[77] http://www.ftc.gov/bcp/conline/pubs/buspubs/canspam.htm
[78] *CAN-SPAM Act of 2003*. Section 7(f)(3)(B), 7(g)(3)(B). Note that the "uncapped" damages are for violations of section 5(a)(1).

hardest to slow by technological means. Enforcement shifts the costs and burdens of spam back to spammers themselves, where they belong.

It is important to realize that the CAN-SPAM Act is only the latest among the federal and state laws that can be used against spammers. For example, in May 2004 the notorious "Buffalo spammer," believed to be responsible for over 825 million spam messages, was sentenced under New York state forgery and identity theft laws to 3 to 7 years in prison.[79] Interagency and public-private collaboration is also considered key in the fight against spam. The FBI testified in May 2004 that as part of its "SLAM-Spam" initiative it sought to increase collaboration with industry, the FTC, and State Attorneys General.[80] Nonetheless, according to Sen. John McCain, five months after CAN-SPAM's passage, government and industry together had brought only eight cases against spammers.[81]

The major enforcement problem is not proving that a spammer has spammed but tracking down and identifying the spammer. Enforcement officials are loath to publicize detailed accounts of their techniques but broadly speaking, they are similar to those used in locating the senders of any fraudulent, deceptive commercial communications. For example, the FTCfollows communications, the money transferred, and the path of the goods delivered to trace back and find the sender.[82]

Spammers may use forgery to interfere with each of those three routes, the easiest being communications. Introducing an international intermediary into the chain can be a

---

[79] *"Buffalo Spammer" Sent to Slammer*. Reuters. May 27, 2004.
[80] Jana D. Monroe. *Testimony to US Senate Committee on Commerce, Science, and Transportation*. Federal Bureau of Investigation. May 20, 2004.

[81] John McCain. *Testimony to US Senate Committee on Commerce, Science, and Transportation*. US Senate. May 20, 2004.
[82] Personal interview, Dan Salsburg. July 16, 2004.

very effective way to create problems for investigators (discussed in an upcoming section). In testimony before the FTC, Paula Sellis, an attorney with the Washington State Attorney General's office, described a spam prosecution that required 14 pre-suit subpoenas to ISPs and banks. The process took months. [83]

Anti-spam enforcement to date has not been a complete failure. In March 2004, a group of large ISPs announced six joint lawsuits under the CAN-SPAM Act against hundred of alleged spammers, many of them "John Doe" defendants.[84] In April 2004, America Online gave away a Porsche Boxster, which it had seized from a spammer as part of a legal battle a year earlier.[85] In May 2004, Virginia's Attorney General announced the indictment of a spammer for allegedly violating Virginia's state anti-spam law.[86]

The real value of such enforcement is not the positive press it generates for those taking a hard stance against spam, but the volume of spam it can stop as a deterrent. There is no easy way to quantify that reduction. ISPs, state governments and the federal government have assigned considerable resources to this fight and it is sometimes possible to win substantial civil judgments from spammers. Given that the CAN-SPAM Act also allows for recovery of attorneys' fees (for states and ISPs),[87] enforcement can be more than symbolic.

In fact, there have not yet been any *convictions or judgments* pursuant to CAN-SPAM. The strict regulations and requirements set forth in the law are meaningless if they are not followed. The stiff penalties are irrelevant if they are not levied against

[83] *Do Not E-Mail Registry Meeting*. Federal Trade Commission. March 10, 2004. 15.
[84] Press Release, Microsoft. March 10, 2004.
[85] Press Release, American Online. April 29, 2004.
[86] Press Release, Virginia Attorney General. May 25, 2004.
[87] *CAN-SPAM Act of 2003*. Section 7(f)(4), 7(g)(4).

violators. Infinitesimal compliance rates and a growing volume of spam strengthen the conclusion that most spammers will not change their practices just because they are now illegal. A poorly-enforced anti-spam law is almost as bad as no law at all.

Individual Action and Bounties

Individuals and businesses have no direct recourse under CAN-SPAM. Such a provision may have been omitted in response to a Utah spam statute (later repealed) that provided for individual lawsuits against spammers and "resulted in thousands of class-action lawsuits against legitimate marketers and Internet service providers who did not send spam."[88] Instead, the Act calls on the FTC to study the efficacy of a bounty system, under which individuals who supply information resulting in the collection of a civil penalty from a spammer would be entitled to not less than 20 percent of those funds.[89]

In September 2002, law professor Lawrence Lessig suggested that individuals who provided proof to the FTC that spammers had violated a spam law be entitled to a monetary reward.[90] Lessig saw this as a way to leverage the larger Internet community's knowledge and antagonism toward spammers (as illustrated, for example, by the Register Of Known Spam Operations, or ROSKO[91]) without resorting to what he perceived as overzealousness in the operation of certain IP blacklisting organizations.

Today, given that the FTC has filed only two anti-spam lawsuits under CAN-SPAM,[92] the effectiveness of bounties in encouraging enforcement appears questionable. On the other hand, individual action appears more promising because it would almost

---

[88] Unspam
[89] *CAN-SPAM Act of 2003*. Section 11(1).
[90] Lawrence Lessig. *Code Breaking: A Bounty on Spammers*. CIO Insight. September 22, 2002.
[91] http://www.spamhaus.org/rokso/
[92] Timothy Murris. *Testimony to US Senate Committee on Commerce, Science, and Transportation*. Federal Trade Commission. May 20, 2004. The FTC has no press releases about CAN-SPAM lawsuits after this date.

surely lead to the filing of more cases against spammers. The most significant drawback to allowing individual action is the risk of frivolous or profiteering lawsuits against innocent non-spammers (as in Utah) and the large level of both e-mail volume and technical and forensic expertise that would be required for an individual to bring suit. Note that large (non-ISP) technology companies might be well suited to such a task.

If the entities that the CAN-SPAM Act has authorized to take action are unable to enforce the law effectively, then it makes sense, with due caution and care, to empower individuals to take matters into their own hands.

State Preemption

Another mixed outcome of the CAN-SPAM Act was that it superseded all State spam laws except to the extent that such a law "prohibits falsity or deception in any portion of a commercial electronic mail message."[93] This preempted state laws like California's, which required marketers to adhere to opt-in standards.[94] In April 2003, the Internet Committee of the National Association of Attorneys General wrote to members of Congress complaining about the law. "Twenty-seven states have enacted laws targeting unsolicited commercial e-mail… the protections in the proposed legislation fail to approach all those in state laws. We are troubled that the proposed legislation would preempt states from enacting and enforcing anti-spam provisions that exceed the scope of federal legislation."[95]

While there are clear advantages to a single national law to which all marketers must comply, some state legislators and Attorneys General are disappointed that the

[93] *CAN-SPAM Act of 2003*. Section 8(b)(1).
[94] *California Senate Bill 186*. 2003.
[95] Letter, National Association of Attorneys General. April 29, 2003. 1.

CAN-SPAM Act has made it easier for spammers to legally reach the residents of their states.

Judging Success

Determining how to evaluate the success of the CAN-SPAM Act poses its own problems. Many factors contribute to the growth of spam and the causal relationships between them are not always clear. If the bellwether spam percentage numbers fall, that could be the result of the Act, deployment of new anti-spam technology or increased international cooperation. Of course, anti-spam vendors, who are responsible for the vast majority of such data, do have their own agendas. Lawsuits under the Act provide anecdotal evidence that it is a useful tool, but it is by no means certain that they reduce spam to any meaningful extent. Further, CAN-SPAM could succeed even without evidence of spam reduction. Volume analysis, consumer reaction, and even compliance figures may not be able to capture a qualitative change in the nature of spam, which is, arguably, a goal of the CAN-SPAM Act. The lack of a clear measure of success from the start means that the government has little direction other than to devote resources and attack the problem as best it can on an ad-hoc basis. The FTC effectiveness report will probably be the first real world, systematic and public examination of the Act's real impacts.

Recommendations

Given the current outcry over spam and the consumer demand for solutions, technology (i.e., the market) will eventually win the war against spam. However, without a deterrent and lacking the ability to seek recourse provided by strong anti-spam laws and

their enforcement, the costs of this victory will lie squarely on the shoulders of e-mail users. It will likely take longer than it would under a system of strong laws and enforcement.

To address the threats outlined in the introduction, policy makers should:

- **Ensure that the CAN-SPAM Act is enforced vigorously**

Enforcement under the Act has been mixed. So far, ISPs have been responsible for the majority of lawsuits. The FTC and the Department of Justice are the federal agencies charged with enforcing the CAN-SPAM Act but have yet to pursue large numbers of spammers under the Act. The overwhelming majority of spam today violates the CAN-SPAM Act and a large amount of that is flagrantly illegal. More lawsuits, and more victories, are critical.

- **Create an individual right of action**

One way to almost certainly increase the volume of litigation against spammers is to allow individuals to sue them directly. This is probably the most risky recommendation because poorly worded legislation could encourage frivolous litigation. Also, allowing individual action would likely require further legislation, which presents an obstacle. Nonetheless, granting individual action is probably the best way to ensure that spammers are held accountable for their actions.

- **Establish oversight for the technology process**

The federal government does not have the private sector's level of spam-fighting technical expertise and it would be inappropriate for government to choose or write specific anti-spam technologies into law. However, the government does have two

important interests in the technology process. First, some technologies can help make enforcement easier. Second, the public sector can help to ensure that eventual solutions remain open, rather than closed and proprietary. (In this regard, the FTC is planning on holding a summit on sender authentication technologies in the fall of 2004.)

- **Revisit privacy concerns**

The CAN-SPAM Act does allow the sending of e-mail that many consumers would consider spam. While this issue may be less important today than the fight against unethical outlaw spammers, it is likely to reemerge when significant headway has been made on the latter. At that time, it will be important to consider instituting something like ADV subject line filtering, a Do Not Email Registry, or opt-in.

- **Promote user education**

The public sector may be best suited to address the difficult job of user education. The fewer the number of people who purchase from spammers, fall victim to e-mail based identity theft, or even run misconfigured and vulnerable computer servers, the more difficult it is for spammers to earn a profit. In January 2004, the FTC launched Operation Secure Your Server, an initiative to identify and educate the owners of misconfigured servers, in collaboration with 36 other agencies in 26 countries.