Electronic Frontier Foundation

Defending Freedom in the Digital World

Noncommercial Email Lists: **Collateral Damage in the Fight Against Spam**

By Cindy Cohn and Annalee Newitz

I. The Problem

MoveOn.org is a politically progressive organization that engages in online activism. For the most part, its work consists of sending out action alerts to its members via email lists. Often, these alerts will ask subscribers to send letters to their representatives about time-sensitive issues, or provide details about upcoming political events. Although people on the MoveOn.org email lists have specifically requested to receive these alerts, many large ISPs regularly block them because they assume bulk email is spam. As a result, concerned citizens do not receive timely news about political issues that they want. Often, MoveOn.org's staff doesn't discover that the mail isn't getting through for days or weeks, and even when it does, ISPs respond slowly to "unblock" requests or refuse to explain why email has been confiscated. Although ISPs may have the best of intentions, what we see in this scenario – one that is all too common –is free speech being chilled in the service of blocking spam.

In their zeal to stop spam, many organizations and companies are blocking the delivery of wanted messages, especially those sent through email lists. This problem is exacerbated by the fact that most blocking processes are not transparent to the email sender or recipient, and email users are generally given little or no control over which emails are blocked. Instead, system administrators, creators of spam-blocking tools, and ISPs all too often attempt to predict what mail a recipient does and does not want. As a result, email users rarely receive all legitimate messages sent to them.

The large number of anti-spam tools is a tremendous problem for email list owners, who must navigate everything from “block lists” to "Bayesian filters" to communicate with willing recipients. The fact that unwanted email often masquerades as wanted email complicates matters, as do the ongoing differences of opinion and policy about when a person has consented to be added to an email list. There is also some evidence that administrators are misusing spam blockers to block email lists because of personal malice or political opposition to the content of the messages. This is clearly the case when email is administered under government regimes like the one in China.

Additionally, a growing number of proposals, loosely called “bonded sender” initiatives, require organizations sending bulk email to pay a fee to register with various "bonder" organizations. This practice might mean that groups that cannot pay will have their noncommercial email relegated to second-class status. Indeed, expensive certification requirements and reflexive blocking of all "uncertified" email could mean that mail from noncommercial mailing lists won’t be delivered at all.

When tools designed to prevent unwanted email also prevent wanted email from being delivered, or when anti-spam tools favor well-funded speakers over others, something fundamental to the health of Internet communication has been broken. Email is no longer a strong vehicle for free speech.

In an effort to resolve this problem, a coalition of noncommercial email mailing list owners has joined with the Electronic Frontier Foundation (EFF) to create a list of principles[1] for mailing list owners, ISPs, and anti-spam forces. This guide aims to help all three work together to distinguish better between wanted and unwanted mailing list messages.[2] Our goal is to ensure that Internet users receive all of the email that they want to receive, without being inundated with unwanted messages. At the same time, we want to preserve the ability to send bulk email to lists of people who have chosen to receive it -- something spam-blocking technologies and policies threaten to burden, if not eliminate . In this paper, we introduce the major problems faced by senders and receivers of noncommercial bulk email, offering suggestions for best principles and practices in spam management.

Noncommercial Email Lists

Email lists are among the most important, powerful, and accessible communication tools on the Internet, allowing a single person or group to send messages to a much larger group of people who have agreed to

[1] This document does not attempt to outline or meet the full legal requirements for ISPs, mailing list owners, or anti-spam tools. Thus, the document suggests that noncommercial mailing list owners receive affirmative confirmation before adding a person onto the list, even though the First Amendment would likely prevent institution of this as a legal requirement. Similarly, while there are possible legal theories, there is no clear legal rule that requires an ISP to deliver all wanted email to its customers, or that punishes spam filters for failing to correct errors. The goal of this document is to suggest a reasonable way forward for all concerned, not to advocate or rule out any legal tests, duties, or restrictions.

[2] The focus of this effort is noncommercial email mailing lists, because they are a key locus of free speech and usually have fewer resources to discover, track, and follow up when emails are blocked. Commercial list serves and commercial and noncommercial one-time emails are outside the scope of this document.

receive the messages. They allow recipients to learn about current issues and participate more easily in initiatives and events that they care about. The topics addressed by noncommercial email lists are as diverse as human thought itself; there are lists devoted to electoral politics, AIDS prevention, knitting, and the San Francisco 49ers. Many lists exist much more informally, uniting groups of friends or colleagues in ways that defy categorization.

As email technology has matured, individuals have started to rely on mailing lists for critical information. For example, courts nationwide give notifications about ongoing litigation via email lists for counsel in a case. Government entities ranging from taxing authorities to business registration agencies are using email for notification and processing of critical information, again often using mailing lists. Email lists give people the ability to track government and world events minute-by-minute, and thereby participate in public debate in new and powerful ways.

Yet the continued viability of email lists as a cheap, efficient means of one-to-many communication is at risk. An informal survey conducted by EFF in 2003 revealed that many organizations with large email lists, and even some organizations with smaller ones, face an ongoing struggle to get email delivered to members. List owners for groups as small as the parents of Berkeley, California high school students and as large as Moveon.org, which has lists with two million subscribers, reported problems with anti-spam mechanisms. Other list owners negatively affected by these mechanisms include technologist and author Bruce Schneier, who publishes the highly respected Cryptogram newsletter, and the people behind TidBITS, a prominent email list for the Macintosh Internet community. EFF faces ongoing difficulties with anti-spam mechanisms in sending out our own long-running newsletter, EFFector.

There are multiple issues email list owners and recipients face as a result of to these mechanisms. Below is a list of some of the major problems, although it is by no means comprehensive.

Lack of Transparency

It's difficult for list members to figure out that their email isn't being delivered. Recipients report that they don't notice that they've stopped receiving messages for several weeks or months, and often only after missing important ones. Similarly, email list owners say that it's hard to know when their messages have been blocked. Often, they only discover blocks when they receive angry or confused messages from subscribers who believe they've been dropped off a list intentionally or through negligence. Some blocks result in bounced messages to the

email list owner, providing an explanation of what went wrong -- but most blocks do not. And no email list owner or recipient is warned ahead of time that a message will be blocked, much less receives instructions about how to avoid it.

Even when an email-list member discovers that her mail is being blocked, it's often extremely difficult to find out who has blocked it and why. While her ISP is usually the direct cause of the block, ISPs generally use a software package or third-party anti-spam service such as MAPS or SpamCop,[3] and that list or mechanism is what determines whether or not a message is delivered. Tracking down the proprietors of blocking software and anti-spam services can be very difficult. ISPs are not usually forthcoming with the names of the various private services they use, even to subscribers, and anti-spam services rarely list their clients. Moreover, an anti-spam service often won't reveal its rationale for blocking certain senders, even to the ISPs with whom it does business. Thus, even if an ISP admin wanted to explain to a user why he hasn't received his email, often she can't.

But suppose that an email list member bypasses the ISP and ferrets out who or what has blocked his mail. Unfortunately, the reason for the block is even more likely to remain a mystery. Many in the anti-spam community prefer to use secret rules and algorithms to decide what messages will or will not be delivered. Some do so for competitive reasons, others for strategic purposes. Regardless, the end result is that the spam solution provider is not likely to provide the individual with the "magic recipe." Thus, there is often no way to avoid the block by learning the rules in advance, and given the fact that recipes change over time, doing "forensic analysis" after the fact is frequently a fruitless endeavor.

Common Problems with the Rationales for Blocking

Email is typically blocked for a few basic reasons, some more fair than others. Here we outline four techniques that inform email blocks, and that can lead to situations where people aren't getting the emails they wish to receive.

Probabilistic Classification/Machine Learning

Probabilistic classification is a family of techniques involving computer programs that "learn" what is and is not spam, allowing the programs to adapt over time. Using "machine learning" algorithms, the programs determine the probability that a given email should be classified as spam. The technique

[3] For an explanation of how blocklists like MAPS work, see <http://www.scconsult.com/bill/dnsblhelp.html>.

known as "Bayesian filtering" belongs to this group.

These algorithms must be “trained” with starter data before they can begin automatically classifying documents. Different learning algorithms achieve varying degrees of success, but most such algorithms improve as they are trained by users who mark certain mail as "spam" and other mail as "not spam." As a result, this technique can allow for significant end-user control.

Ad Hoc Pattern Matching

Many spam detectors search for specific spam-like patterns, such as all-caps or gappy text, words like “Viagra” and “mortgage,” misspellings, strings of numbers in the subject line header, non-Latin character sets, and the like. The exact patterns used vary widely and are in constant flux. Some people in the anti-spam community take the position that the use of certain words is equivalent to sending spam, regardless of the fact that these words have legitimate uses. EFF has often been a victim of this: we have been told that EFF's email newsletter will not be delivered unless we stop using words like "spam," "pornography," and "opt-in." One EFF newsletter was blocked as spam because it referred to a group called “Stop Prisoner Rape.” While it's unlikely that EFF's messages are themselves the intended target of anti-spam mechanisms, they are nonetheless blocked due to these imprecise, overbroad techniques.

Spam Assassin, a popular program that does ad hoc pattern matching, assigns "points" to various features of an email to determine whether it is spam. The higher the number of points, the more likely it will be sent to the spam folder or discarded. Points can be assigned for everything from country of origin to certain words or subject headers. One of the major problems with this system is that messages from certain countries – like China, for example – can be blocked purely on the basis of where they come from and what language they’re in The implications for free speech here are very troubling indeed: a human rights group communicating with people in China may find that their bulk email is blocked, and thus anti-spam technology unintentionally works as a political censorship mechanism. Of course, this is only a problem when end users are not given control over how points are assigned, and what will be done with messages that get "high" or "low" marks. Spam Assassin and programs like it can be configured to give users more control.

Collaborative Classification

With this system, users classify documents as spam or not spam, and this classification is sent to a central server. When new lists of classifications are sent to the server, it checks to see whether or not other users have classified the same messages as spam. Thus, a community of people can work together to filter spam. Vipul Ved Prakash's Razor system, as well as the Distributed Checksum Clearinghouse (DCC), work this way. Like Spam Assassin, this technique has the advantage that it can be deployed in a way that gives control to end users.

Blocklists and Whitelists

In this method, some self-appointed authority compiles block lists (and occasionally, white lists) of domain names and/or IP addresses, then publishes the lists on the Internet. Email server operators can subscribe to the block list service and instruct their email servers to deny receipt of email from the listed hosts. This is generally not something the end user has control over, since a key purpose is to block spam at the SMTP interface, thereby saving bandwidth.

A common form of block list is a list of IP addresses to block, including one or more hosts alleged to have sent spam. The express purpose of this technique is to cause collateral damage, forcibly involving more people in the block list compiler's "cause." Transplanted "offline," this kind of policy would hold that it's reasonable to boycott a store that uses a specific long-distance telephone company simply because the telephone company (not the store) also provides long distance service to someone you dislike. A policy like this is clearly unjust to non-spamming hosts, given that it subjects them to poor treatment simply because they share an ISP with an alleged bad actor.

Occasionally, block lists will block all dynamic and dialup IP ranges, despite the fact that these IP ranges have perfectly legitimate uses. This practice also makes it difficult for tiny nonprofit organizations to set up their own mail servers.

In addition, some sites are added to a block list because of the procedures followed by the operators of the email servers at the site; for example, email servers that are configured as open relays (meaning anyone can use them to send email to anyone). The justification for this is that spammers use such servers to hide their identities, despite the fact that open relays have legitimate non-spam uses.

Email Authentication

Email authentication technologies are intended to help positively identify the server sending a message, and are supposed to cut down on spam messages that "spoof" the identity of sending servers. The idea is to stop people from using fake email addresses to send spam.

Typical systems that enable email authentication include Sender Policy Framework (SPF), SenderID and DomainKeys. These methods enable recipients to confirm that email is from the domain it appears to be from. All three systems share a reliance on augmentations to the Domain Name System (DNS), which links IP addresses to domain names. DNS records have been expanded so that domain owners can identify the specific mail servers authorized to send mail for their domain. When you receive mail purporting to be from Example.net domain, your server might use sender authentication to see if the sending mail server is authorized to send mail from Example.net. Most groups using sender authentication say that if an email fails the authentication test, it

is a strong indication that the mail has a forged sender and probably should be blocked.

SPF, SenderID and DomainKeys differ in the specific component of an email message that each tests. SPF (which was recently adopted by AOL) is simplest – it checks the "envelope sender" of an email (which includes the domain name of the mail server initiating an SMTP connection). SenderID delays its checking until after message data are transmitted, and examines several sender-related fields in the headers of an email message to identify the "purported responsible address." DomainKeys checks a header containing a digital signature of the message body and certain parts of the header. This system is more complicated because it verifies the domain of each email sender (the actual "from" address a recipient sees) as well as the integrity of the message.

Many have described the email authentication systems as promoting a policy that says email is “spam unless proven otherwise.”

Antispam policies based on email authentication can also hinder free speech, as activists participating online letter-writing campaigns have discovered. The software that enables activist letter-writing campaigns on the net is designed to make it easier for concerned citizens to write email to their representatives about pressing political issues. A concerned citizen writes her letter in an online form and indicates in a checkbox which representative or public official she wishes to reach. The activist campaign software then sends the email on her behalf, putting the letter-writer's email address in the "from" field but sending it from servers at the activist organization providing the service. Unfortunately, emails sent in this fashion appear "spoofed" to email authentication software because the sender's domain is different from the domain where the email originates. One activist reported to the EFF that when she used letter-writing campaign software to tell her senator how she felt about some upcoming legislation, her emails were turned away because he had used SPF email authentication on his server.

Lack of Due Process

The RFC standard for SMTP email protocol defines a duty to deliver or report back on non-delivery.[4] Yet increasingly anti-spam mechanisms and the ISPs that use them are deviating from this requirement in cases of suspected spam. This is unfortunate, as the RFC serves a real purpose – to keep email flowing and to assist in the detection and correction of errors.

Outside of their RFC duties, spam blockers and ISPs generally have no specific legal obligation to provide any sort of due process when they

[4] "The responsibility of an SMTP client is to transfer mail messages to one or more SMTP servers, or report its failure to do so."
<http://www.networksorcery.com/enp/rfc/rfc2821.txt section 2.1>

choose to block a message, sender, or an entire IP block. They also have no specific legal obligation to ensure that these blocks are removed when they have been wrongly implemented, or when the spamming ceases. Anecdotal reports indicate that some anti-spam services take up to two weeks or longer to remove a sender from a block list. Others report that no process exists at all. Some even claim that anti-spam services are charging senders a fee to be removed from the block list. Obviously these policies create tremendous opportunities for misuse, especially when no objective criteria or requirements for blocking or unblocking exist.

Delays in getting a sender removed from a block list have a huge impact on political organizations attempting to provide timely information. MoveOn.org uses its list is to give recipients up-to-the-minute information about breaking news and political and cultural events. Recipients rely on the list to do things like help them write their elected representatives before the deadline for a vote on a specific bill before Congress. Since these deadlines are critical as the time for decision approaches, even a slight delay can effectively prevent an email list recipient from making his or her voice heard in the democratic process.

Anti-Competitive, Spiteful, and Politically Motivated Blocking

There is increasing concern that anti-spam measures are being misused. There are a number of reports suggesting that individuals and groups have been labeled as spammers out of personal malice, anti-competitive behavior, or even a fits of pique. –For example, the technology journalist Declan McCullagh reports that SpamCop blacklisted his email list, Politech, evidently because of an alleged spammer's personal vendetta against him. McCullagh had flagged the individual as a spammer by emailing abuse@yahoo.com, so the accused spammer reportedly sought revenge by likewise reporting McCullagh to SpamCop. Without checking on the source of the report, SpamCop listed McCullagh as a spammer. Rectifying the situation proved difficult, and McCullagh was incorrectly listed as a spammer with SpamCop two more times after that.[5]

Bonded Senders: Barriers to Entry

A number of ISPs and companies like IronPort and TrustE have begun developing "bonded sender" programs. While details vary, the

[5] See <http://www.politechbot.com/p-03730.html>, <http://www.politechbot.com/p-04121.html>, <http://www.politechbot.com/p-03372.html>.

basic premise of these programs is that only entities or persons who have been “certified” will get their email list messages delivered in a timely or prioritized manner (or, taken to its extreme, delivered at all). Essentially, these programs empower certain entities and organizations to serve as gatekeepers for bulk Internet mail.

These mechanisms are troubling because they could lead to a situation where small players (in terms of funding rather than size of the recipient list) will be unable to use email lists to reach subscribers. Worse, since these mechanisms dock a monetary "bond" whenever the bonded sender companies receive a certain number of spam complaints, they create a situation ripe for manipulation by political enemies or competitors. If someone doesn't like a particular group's message, he or she can report the group as a spammer and actually cost the group money. This wouldn't be a problem except for the fact that most bonded sender programs have no way to check the authenticity of complaints against a given mailer. Moreover, some have acknowledged that they have no formal plans or processes to do so. False or politically-motivated complaints will punish legitimate mailers as if they were spammers.

Another problem with bonded sender programs is that they push email into becoming a “pay to play" medium, where people with money can eat their fines and have email delivered on a priority basis, while those who with less money face unreliable delivery. While paying to get email prioritized is not a new development online, the “bonded sender” programs would worsen the problem, perhaps resulting in a world where organizations without financial resources or connections will get their email delivered late or not at all.

Conclusion

Anti-spam measures can and should be deployed as part of email systems. But those who implement these measures must be sensitive to the fact that what they are processing is speech, and that free speech is one of the core elements of a democratic society. If anti-spam measures are preventing wanted speech from reaching a willing recipient, whether intentionally or unwittingly, they are hurting free speech. If they create additional costs or red tape for groups sending noncommercial bulk email, they are damaging one of the core benefits of the Internet: the level playing field for speakers.

II. The Solution (Or At Least A Start): Principles and Best Practices

EFF has developed four basic principles and several best practices that should be applied to all email traffic. These are based upon the fundamental ideas of

user control, information transparency, and fair play.

1. Individual recipients should have ultimate control over whether they receive the messages they wish to receive. They can be assisted by software or anti-spam services, but knowledge of and control over receipt of email should remain with recipients and end users.

2. All mailing-list email should be delivered to willing subscribers. As a corollary, no one should be subscribed to an email list without his or her knowledge and consent, as evidenced by positive action.

3. Developers and proprietors of anti-spam technologies should avoid solutions that are overbroad and easily abused, and ISPs should likewise avoid implementing solutions that are overbroad and easily abused.

4. Anti-spam measures must be sufficiently transparent to allow email list senders and recipients to discover when and why their email is being blocked in a timely fashion. This means that anti-spam services and abuse departments must respond to user queries quickly.

Best Practices for Email List Owners

1. Senders must ensure that recipients have taken positive action indicating that they wish to be signed up for a mailing list. While this problem is less of an issue with noncommercial lists, recipients do report that they have been added to noncommercial mailing lists without their consent. Sometimes this happens after they participated in a single call-to-action or responded to an issue online. Other times, organizers use or purchase a mailing list set up for one purpose as a "starter list" for another, with the incorrect assumption that the people on the first list are likely to be interested in the second. Occasionally, people will be added to email lists by someone spoofing their email address and requesting signup.

Senders also have a duty to ensure that the process by which recipients join a list is transparent. The recipient must take some sort of positive action to indicate that he or she has consented to receive mail from each list.

Positive action can take any number of forms, but clearly includes:

- Checking a box or otherwise affirmatively marking a web

form with no pre-checked boxes, as long as this action is followed by a confirmed opt-in email [6]

- Sending an affirmative email to the mailing list owner (since affirmation emails can be spoofed, this positive action must be used in concert with others)

- Signing up offline via petition or other mechanism, as long as the offline form contains clear notice that the signer will be added to an email list

Positive action does *not* include:

- Submitting pre-checked, default-subscribe forms

- Submitting information when what the user is agreeing to is not visible on the same screen where the positive action is taken. (A common example of this sort of unacceptable technique is a link to a privacy policy, where, often in opaque legal language, a policy of sharing broadly with affiliates is revealed.)

- Using information gathered offline for other purposes – for example, a sign-in sheet at a meeting that fails to reveal that it will be used to create a mailing list.

2. Senders should provide ongoing information about how to unsubscribe. In addition to positive action to add an email address to a mailing list, the first message from an email list should ask the user to confirm that he or she has indeed requested to join the mailing list, and contain a link to a site where the user can confirm his or her desire to be on the list. All messages thereafter must plainly and clearly inform the user that she has signed up for a mailing list, as well as provide easy-to-use instructions to unsubscribe. Ideally, users should be able to unsubscribe with a single action, like a web link or reply email – unsubscribing should not require a complicated login, provision of a password, etc. Senders must always make unsubscribe information easily available – possibly via information on the bottom of each email sent [7] -- and respond expeditiously to unsubscribe requests.

---

[6] For more information about confirmed opt-in, see http://www.aota.net/Mailing_Lists/Confirmed-Opt-In.php4

[7] Currently the problem is that senders of unwanted email use these systems to harvest true email addresses (as opposed to many false ones). However, if the initial welcome message comes from a known sender, the risk to the subscriber in responding is lower.

3. Senders must engage in "list cleanliness." When an email list has out-of-date email addresses and fails to remove them from the list, this causes a burden for the receiving ISP. Email list owners should remove an email address from the list upon receipt of one "hard" bounce (a response from a recipient's mail server saying there is no such user), or three "soft" bounces (responses indicating that the email address is unavailable for some reason).

Best Practices for ISPs and Anti-Spam Services

1. ISPs and anti-spam services should ensure that recipients have ultimate control over any anti-spam mechanism that affects the email they receive. Most anti-spam mechanisms work in one of two ways: blocking or filtering. In spam-blocking systems, certain email messages aren't delivered and the recipient never receives any indication that they were sent. In spam filtering, email is screened prior to delivery and either 1) scored for spam probability so that the recipient can quickly identify probable spam, or 2) automatically quarantined in special mailboxes such as a bulk mail folder or spam folder.

Best practices favor filtering over blocking. While there are exceptions, filtering or routing is generally preferable to blocking suspected spam because it gives users the opportunity to correct an improperly labeled email. Blocking based upon user input – user-created block lists applied only to that user – is acceptable. Exceptions to this suggested practice can include blocking email at the server level that has been determined to be spam through a "honeypot" system[8] or some other system that does not rely on the content of the messages or pre-approval of senders.

Filtering methods must empower the recipient. The best method for ensuring that wanted mail is delivered is to place the tools in the hands of the recipient, on the client side. This does not mean that the recipient sees every email, but rather that recipients can turn off filtering, configure it, train it, review blocked messages and unblock (and block as desired) specific senders and categories of senders.

2. Senders and recipients should be able to determine when emails have been blocked and should be told why. The failure to inform either a sender or a recipient when an intermediary has prevented delivery of email creates a number of problems, some of which

[8] For a complete explanation of spam honeypots, see Laurent Oudot's article in Security Focus at http://www.securityfocus.com./infocus/1747.

actually exacerbate the problem of unwanted email. Not only does it prevent recipients from learning when their messages have been incorrectly blocked by their ISP, it also prevents senders from cleaning up their email lists to remove outdated addresses.

Best practices would include implementing a method that allows both senders and recipients to learn when email has not been delivered. ISPs could post the IP addresses of senders whose messages have been blocked on a website accessible to senders and recipients. They could also notify senders when a recipient has marked their messages as spam. A one-for-one notification when delivery fails might not be reasonable in all instances. But ISPs should have a method for conveying this information or making it available to the sender and recipient.

3. Recipients should be notified if they are unsubscribed from a mailing list and informed about the circumstances that brought this about. Some email list owners report that they receive demands from ISPs that they unsubscribe users with a corresponding demand that they not notify the subscriber that this is occurring or why. List owners should always quickly unsubscribe individuals who no longer want to receive their messages, but the unsubscribe process should be controlled by the recipient, not the recipient's ISP or another third party.

4. Anti-spam services and ISPs must judge email and senders on their own merits, not the actions of others. The practice of blocking or filtering email from innocent senders based upon such factors as the web services they use should cease. Email should not be blocked based upon "bad" IP addresses if the addresses are also used by legitimate groups for sending legitimate email.

5. Anti-spam services and ISPs should cease using blind keyword or phrase blocking. Blind keyword or phrase blocking is the determination that messages will not be delivered because they contain specific words or phrases. This method is imprecise and unnecessary, especially now that more sophisticated tools are available. Moreover, it can be used to block messages for political reasons. In short, there's no defensible reason to label email as spam based solely on keywords or phrases.

6. Anti-spam techniques must allow for quick correction. All anti-spam techniques represent an attempt to use shortcuts to determine whether a particular message is wanted or unwanted by the recipient. Invariably, these shortcuts are imprecise. Since this is the case, there should be a ready method to correct the inevitable

mistakes. Corrections can be made using technology or by personal intervention, but the method should be readily accessible to senders and recipients. This applies to both direct anti-spam technologies and systems such as bonded sender.

7. Anti-spam techniques must not be easily misused. ISPs and anti-spam services should ensure that anti-spam techniques are used in good faith, in an objectively fair and nonpartisan manner. Techniques that are misused or that are easily misused, such as the automatic blocking of all email from a sender if a small number of recipients complain, should not be implemented. People using these techniques have succeeded in singling out and "censoring" politically controversial email message.

8. "Bonded sender" systems must keep barriers to entry low for groups sending noncommercial email and preserve recipient control. We strongly resist any bonded sender system that requires noncommercial senders to get permission from another body in order to guarantee that mail will be delivered. We resist even more emphatically systems that require senders to pay a fee in order to be considered a legitimate sender of bulk email or that charge money based upon complaints without legitimate mechanisms to investigate complaints and provide due process. However, many major ISPs, including Hotmail and MSN, are using bonded sender programs as part of their spam-management systems. In cases like these, any noncommercial bulk mailer should be guaranteed that its email would be delivered without having to pay for the privilege.

Conclusion

While we endorse fighting spam, we believe strongly that free speech must not fall victim to over-broad, ineffective filtering and blocking.

The goal of the principles and best practices outlined in this paper is to give mailing list owners, ISPs, and anti-spam services some basic guidelines for working together to ensure that noncommercial email lists remain a vital part of the Internet. If an anti-spam technology or service does not abide by these principles, an ISP should not implement it. If a sender does not abide by these principles, it should have no quarrel when its messages are blocked by spam filters.

Together, we can ensure that the effort to fight spam does not chill free speech. Reasonable anti-spam practices will allow mass-distributed, noncommercial speech to flourish.

Joining EFF in this paper and best practice recommendations are:

- MoveOn.org
- Cryptogram
- PoliTech
- Berkeley Parents Network
- GetActive.com
- TidBITS
- Insecure.org and Seclists.org