

realtimepublishers.com<sup>tm</sup>

# *The Definitive Guide<sup>tm</sup> To*

# Email Management and Security

 **singlefin**  
e-mail protection services

*Kevin Beaver*

Chapter 3: Understanding and Preventing Spam .....	44
What's the Big Deal about Spam? .....	45
Scary Spam Statistics .....	46
Example Estimated Cost of Spam.....	46
Additional Spam Costs .....	47
Unintended Side Effects .....	47
Security Implications of Spam.....	48
What's in it for Spammers? .....	48
Spam Laws.....	49
State Anti-Spam Laws .....	50
Tricks Spammers Use .....	51
Identifying Spam.....	52
Tracking Spam Sources .....	53
Sending Through Open Relays .....	54
Additional Spam Propagation Techniques.....	57
Spam Elimination Techniques .....	57
Disable Open Email Relays .....	58
Filtering Methods.....	58
Blacklist Filtering.....	59
Whitelist Filtering .....	60
Signature Filtering .....	61
Heuristics Filtering.....	61
Email Header Filtering.....	62
Content Filtering .....	62
Bayesian Filtering .....	63
Additional Ways to Fight Spam.....	64
Layered Anti-Spam Defenses .....	66
What to Look for in Anti-Spam Solutions.....	68
Summary .....	70

## Copyright Statement

© 2003 Realtimedpublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimedpublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimedpublishers.com, Inc or its web site sponsors. In no event shall Realtimedpublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.


Realtimedpublishers.com and the Realtimedpublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimedpublishers.com, please contact us via e-mail at [info@realtimedpublishers.com](mailto:info@realtimedpublishers.com).

## Chapter 3: Understanding and Preventing Spam

**Spam:** *noun.* 1. Commercial or other junk email sent without the recipient's permission. 2. Spiced, canned ham made by Hormel dating back to the 1930s.

Spam is a word with two very different meanings. For the purposes of this chapter, I'll spare you the pig references and instead talk about spam as it applies to email in our daily business operations. This spam is sometimes referred to as unsolicited commercial email (UCE) or unsolicited bulk email (UBE). In fact, the two words unsolicited and bulk are what define spam. This definition can be expanded to include email jokes, hoaxes, urban legends, and so on.

 The word spam, as it applies to email, has an interesting story. Supposedly, it dates back to various users posting a large number of junk messages on Internet chat rooms and on multi-user dungeon (MUD) message boards. The story goes that people posted the same messages over and over again unnecessarily clogging up the message boards. This senseless blabber was compared to the blabbering waitress in a famous Monty Python's Flying Circus scene where a patron in a restaurant asked what was on the menu. The waitress replied back saying eggs and bacon; egg and spam; egg, bacon, sausage, and spam; spam, egg, spam, spam, bacon; on and on in an annoying fashion. It's amazing how some of our most commonly used words got started!

Although commercial products comprise the majority of spam messages, there seems to be spam for every imaginable subject. As Figure 3.1 illustrates, spam topics range from religion to pornography.

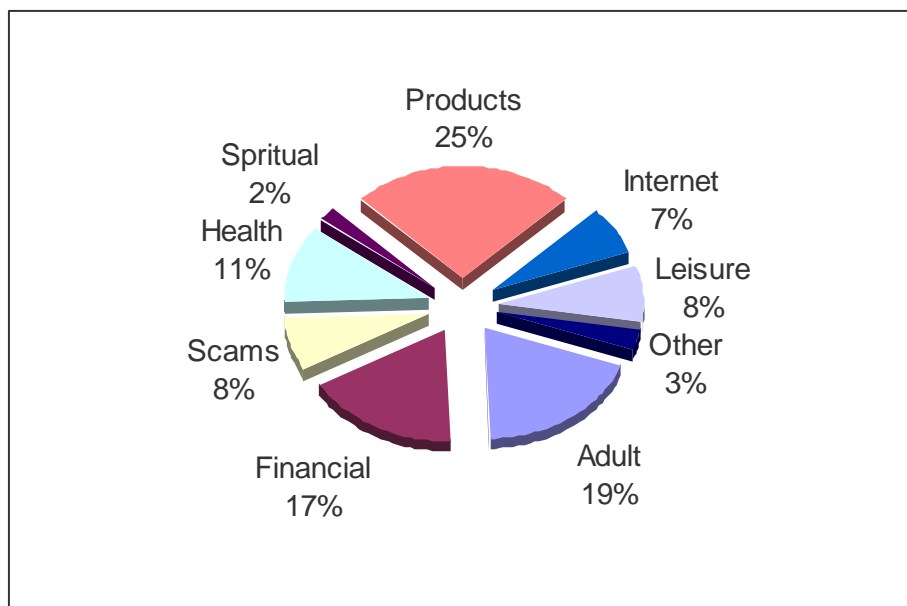



Figure 3.1: Spam category data (Source: Brightmail May 2003).


Did you know that you can lose weight while you sleep, or better yet, save a ton of money on antivirus software for your computer? I've even gotten spam about the latest and greatest no-name spam software. For the purposes of this chapter, I'm going to focus on the messages that are most annoying—the UCE from people you've never heard of trying to sell you something that you probably don't want. These are the messages cluttering your Inbox with subject lines that read like the following (actual spam that I've received):

- Refinance Your Mortgage—Rates at an All Time Low!
- Printer Cartridges—Save Up to 80%—Free Shipping Offer
- Protect Your Computer Against Viruses for Only \$9.95
- Actually reverse aging symptoms
- Save money on your prescriptions
- Approved by many doctors
- Get that girl or guy you always wanted

 I'll discuss other offensive junk email in Chapter 4 in my discussion about content filtering.

There's good news and bad news about spam. The good news is that there are ways to block most of it. The bad news is that the only way to completely rid your network and Inboxes of all spam is to have a human being manually inspect each message as it passes through your systems and only permit the messages that aren't spam. Technically, there's still room for human error with this method, especially as the number of messages increases in the future. I guess another way to completely rid your email of spam would be to stop using email altogether. Although tempting for some of us, neither of these two methods are very good long-term spam fighting options.

So what's required in this growing war we're waging? Enter the world of anti-spam technologies and the policies and procedures that go with them. In this chapter, I'm going to condense what could easily be in an entire book into one chapter. I will cover everything from why spammers do it to how spammers do it to the tips and tricks you can use to prevent spam from infiltrating your email system.

 It's a good time to be in the anti-spam business. According to a study done by IDC, sales for anti-spam products were \$120 million in 2002.


## What's the Big Deal about Spam?

In a nutshell, spam costs organizations time, money, and overall end user productivity. In addition, if spam isn't dealt with properly, your end users may lose confidence in the usefulness of email. Although there is no simple fix for the spam problems we're having, spam cannot be ignored.

## Scary Spam Statistics

Before we delve too far into this chapter, the following list highlights some of my favorite spam statistics that I think will get your attention:

- Spam comprises 55.1 percent of all emails (Source: MessageLabs' May 2003 Monthly Email Security Report)
- Microsoft claims that spam accounts for 80 percent of all Hotmail messages
- 90 percent of all spam received by Internet users in North America and Europe is sent by less than 200 spam outfits (Source: Spamhaus Project)

 To see a list and detailed information about these spam outfits, check out the ROKSO list at <http://www.spamhaus.org/rokso>.


- According to a study performed by the Federal Trade Commission, two-thirds of spam contains false claims, 96 percent of spam offering business and investment opportunities contain false claims, and 48 percent of spam promoting health services or products contains false information.
- One day in early 2003, AOL blocked 1 billion spam messages; its previous high was 780 million blocked spam messages in one day (Source: Direct Newline)
- 4.9 trillion spam messages are projected to be sent in 2003 (Source: Radacati Group)

## Example Estimated Cost of Spam

These statistics justify the war on spam. However, let's look at a real-world example of what spam could actually cost an individual organization. Say the average corporate user receives 50 emails per day (both legitimate emails and spam) Monday thru Friday and another 50 emails over each weekend for a total of 300 emails per week or 15,600 per year. These numbers are fairly conservative, and your spam numbers may vary. (Some reports state that as much as 70 percent or more of email is spam, but I've seen numbers as low as 30 percent.) Let's take a good even number of 50 percent for this example. Given that on average, half of all email is spam, we have a total of 7800 spam messages a year for the average user!

Next, consider how long each user takes to tend to individual spam messages—let's say a very conservative 2 seconds to handle each one; thus, the user consumes 4.33 hours per year dealing with spam! If you conservatively estimate that the average user costs the organization \$40 per hour with salary and benefits, the company is losing \$173.33 per user per year (for the average user).

This amount might seem fairly harmless for smaller organizations that have 10 or so employees, but when you start thinking about organizations that have 100, 1000, or 10,000+ employees, spam costs become a serious problem over time. These numbers add up to \$173,333.33 in 1 year for a 1000-employee organization.

 Looking for a spam cost calculator? There are several on the Internet. Singlefin offers just such a calculator at [http://www.singlefin.net/services/calculator\\_result.php](http://www.singlefin.net/services/calculator_result.php).

Taking this estimate a step further, let's look at the computer hardware that's required to support these kinds of numbers. Let's assume that, based on my non-scientific research, the average spam message is around 5KB in size. Based on the average user receiving 7800 spam messages a year, spam adds up to 38MB of clutter making its way to your email server or the end user's local hard drive every year. Again, 38MB doesn't seem so bad for one or even 10 users, but scale it up to 1000 or more and these organizations now have a serious storage issue on their hands—38GB of storage space for a 1000-user network over 1 year! There's also the issue of backup media space that's required and overall network bandwidth that's being wasted.

### ***Additional Spam Costs***

Although frightening, none of these estimates take into consideration the amount of time and money IT personnel have to spend on:


- Workstation-related problems
- Spam filtering software maintenance
- Increased data management responsibilities
- Email servers running out of storage space
- Email servers having to be rebuilt after a crash
- Purchasing and upgrading workstation and/or server hardware
- Data backups having to be restored

Network administrators have to purchase spam filtering products and more storage space and possibly network bandwidth to handle all the spam. End users end up paying more for larger mailboxes for their Web-based accounts or risk losing legitimate emails. Traveling users dialing up the Internet to check email have to pay more for longer connection times because of increased spam download times. Finally, people are paying more in wireless connection charges for spam they receive via email on mobile devices such as PDAs, cell phones, and so on.

There are dozens of related issues that help justify fighting the spam problem in your organization. The point is that even with conservative estimates, the cost of spam adds up quickly! Ridding yourself of this problem will result in saved money, productivity, sweat equity, computer hardware, and network bandwidth.

### ***Unintended Side Effects***

Anti-spam solutions can actually make email less reliable than it was before. Legitimate emails that are unintentionally blocked can cost organizations money in lost business and intangibles such as reputation and customer loyalty. This unintended side effect of some anti-spam solutions is a critical issue that must be considered when selecting, implementing, and managing anti-spam technologies.

 I'll discuss the pros and cons of the various spam-filtering methods along with some tips about how you can prevent false positives later in this chapter.

## Security Implications of Spam

There are obviously financial, productivity, storage space, and bandwidth losses stemming from spam. However, an often overlooked area when it comes to managing spam is that of information security. Recently, spam seems to be becoming more aggressive with messages that actually try to glean private and confidential information from our networks and computers—a far cry from the typical nuisance spam we’re used to. Like in the malware world, spammers are constantly trying to find ways around existing controls.

Another security concern of spam is that of unwanted spam messages eating up a tremendous amount of storage space. Given enough time and with enough messages coming into a network, the overload can create a DoS condition leading to serious email system downtime. The obvious consequence being email system unavailability on the server and/or client, which can lead to messages not being properly sent or received—and ultimately lost business.

Other spam attacks could be considered network intrusions, especially if they come in with malware attached. Everything from viruses to Trojan horses to Web bugs can wreak havoc on networks, server, and end user systems. In many organizations, any sort of confidential information leakage is intolerable. Perhaps a broader issue related to spam attacks is one of IT and security resource utilization. When IT and security personnel must address spam attacks, they are diverted from other, more important, tasks, which can lead to even more security issues cropping up.

Some spam is actually social engineering at work. Social engineering attacks are notorious for being very malicious yet easily carried out. We can train our end users to not open emails or attachments from people they don’t know, but this training can only go so far. After so much inundation with spam, users tend to end up letting their guard down allowing intrusive messages and malware take its toll.

## What’s in it for Spammers?

So why do spammers send spam? First and foremost, there are high potential payoffs in return for very little effort on the spammer’s part. But you’d think that with the huge majority of people not responding to spam that the spammers would eventually figure out that their way of doing business is not very well received. Given the large numbers of spam messages being sent out, it only takes a very small percentage of spam recipients to reply to and purchase some spam-based offering for the spammers to succeed.



Lawrence Canter and Martha Siegel are two of the original, and perhaps most notorious, spammers on the Internet. These immigration attorneys earned their 15 minutes of fame by posting green card lottery ads on thousands of Usenet groups in 1993. These ads put the Usenet world in an uproar, and the ads were later deleted by a *cancelbot*—an automated message erasing script.

Based on my research, the typical percentage of actual responders and ultimately buyers of spam products is almost always less than 10 percent and is usually much smaller. Even if the rate is only one-tenth of one percent (0.001), if the spammer has sent out a million messages and only get \$2 per sale, that’s still \$20,000 in sales! Until the profits are diminished with the help of increased spam filtering, email-savvy users, and so on, the incentive will remain in place and spammers will continue to send spam.







## Spam Laws

Currently, there are no federal laws in the United States that specifically address spam. Selling spam software or mailing lists is not illegal either. There has been legislation under consideration in the United States for years that would require commercial emailers to identify themselves to their email recipients and offer an easy way for recipients to opt-out of unwanted messages. There have been laws proposed that would establish a nationwide “do not spam” list similar to the new national “do not call” list. However, many countries, including the European Union, have already addressed the spam issue and are doing something about it.


Commercial faxes were outlawed in 1991 via the Telephone Consumer Protection Act because of the time and money losses inflicted upon recipients. Meanwhile, the spam problem keeps getting worse. The spam lobby’s argument is often that it’s their First Amendment right to send spam and that it’s not hurting anyone. They say that consumers around the United States could be deprived of a certain product or service that they wouldn’t have otherwise known about if spam isn’t legal. There are spammers and lobby organizations that have been toying with the definition of spam using semantics and other justifications to say that spam really isn’t spam.

 For an interesting insight into of the politics of spammers, check out the Spamhaus Project article “Spam Definition and Legalization Game” at <http://www.spamhaus.org/newsdog.lasso?article=116>.

Regardless of whether you agree that spam is free speech, it costs people and organizations time and money just like the junk faxes used to (and still do). ISPs, network administrators, and end users all suffer. ISPs have to hire technical experts to deal with spam, and, of course, these expenses are passed on the customer.

 ISPs are one of the biggest victims—and thus opponents—of spam.


Have you received any of the spam messages advertising great deals on antivirus software? I get a few every day. Some of these gray market outfits may be legitimate and some may not. Either way, indications show that, thanks in part to these spammers, we’re headed more toward digital rights management (DRM) technologies from the vendors so that they can better protect their assets. An unfortunate side effect—and one that many of us never consider as resulting from spam—is most likely even more software registration and maintenance headaches than we’re currently experiencing.

 Although there are laws against fraud and other trade violations, if you order from spammers, you have no guarantee that you’ll get your goods or that those goods are not counterfeit.

Conventional methods used to fight spam have not been working effectively. There are current solutions to thwart off spam, however, the global problem still exists and will for a long time to come. Right now, when spammers get caught, the most serious method of discipline that appears to occur is the cancellation of their ISP or co-location account. Solutions range from new technology, industry self-regulation, and legislation. Before the various states enacted spam laws, many lawyers and prosecutors used existing trespassing and forgery laws to fight spammers.

## State Anti-Spam Laws

Currently, 33 states in the United States have anti-spam laws on the books. Local judges are granting multi-million dollar lawsuits against spammers. However, these state spam laws are inconsistent, often “confusing” the spammers because they don’t know where their spam recipients live and thus which laws apply. Thus, the increasing need for federal legislation. The problem with federal laws, or any for that matter, is that spam’s ease of use makes it pretty much impossible to create enforceable spam laws.


 For a listing of the various United State anti-spam state, check out <http://www.spamlaws.com/state>.

Perhaps if end users had a federal law that enabled them to sue spammers, some of the nonsense would stop. The drawback is that these laws wouldn’t affect spammers who move their operations off shore.

There is serious doubt as to whether any legislation will help cut back on spam, much less make it go away. Part of the problem is that a lot of spam comes from outside recipients’ countries. Often laws in one country are not enforced or enforceable in another country. This issue is what will keep federal and other spam laws from being completely effective.

Permission-based emails are working quite well for Internet shopping sites and Web portals. However, that desire for legitimacy and success hasn’t trickled down to the individual spammers. Perhaps more targeted opt-in emails would be a good solution, but most spammers wouldn’t want to go to the trouble. The other possible solution—opt-out—isn’t an ideal solution as ISPs, network administrators, and end users still have to deal with the spam at least once. If a federal opt-out law is ever passed, it would, in effect, legalize spam, possibly making the problem even worse than it is now. Spam would most likely come in through our network borders at an ever greater rate canceling out any planned benefit.

There are various industry self-regulation attempts to thwart the spam problem by vendors such as ePrivacy Group and TRUSTe. These programs attach a “trusted” seal within emails, confirming to the recipient that the email is from an authenticated source. Unfortunately, there’s a downside to this confirmation. The companies that would sign up for such a program are most likely trustworthy companies in the first place. The illegitimate spammers are just going to keep doing what they’ve been doing and not risk getting caught.

 ePrivacy Group recently announced an open standard to help reduce the amount of spam and improve the reliability of legitimate email delivery. The standard is called Trusted Email Open Standard (TEOS), and you can find out more about it at <http://www.eprivacygroup.com/teos>.

There are also organizations such as Habeas’ Sender Warranted Email (<http://www.habeas.com>) that offer email signature services through which legitimate email marketers insert what they call *haiku lines* into their emails signaling that their messages are not spam, so spam filtering methods won’t stop them.

## Tricks Spammers Use

So how do spammers get your email address? There are several ways they glean addresses—some obvious and some you might not have considered. One simple method is to simply manually inspect Web pages. More complex spammers have been known to employ Internet professionals, software developers, and email experts to advance their address-gathering techniques. Perhaps the most common way to access email addresses on the Internet is to use *spambot* software; similar to search engine robots, this software goes from site to site and page to page looking for the easily distinguishable `someone@somedomain.com` combination that denotes an email address. These addresses are gathered from places such as:

- General contact information listing Web sites
- Employee listing Web sites
- Usenet newsgroup postings
- Chat room and other Web-based discussion group postings
- General email discussion groups


An even more basic way that spammers can obtain tens if not hundreds or thousands of email addresses is to obtain forwarded jokes or even legitimate email when the sender does not blind copy (BCC) everyone. Don't you just hate it when you're on the receiving end of one of those privacy breaching emails?

A newer, more clever method of obtaining email addresses is called a directory harvest attack (DHA) in which spammers have automated tools that will send emails to hundreds, if not thousands, of possible email addresses for a specific domain. Those emails that do not bounce back are assumed to be legitimate addresses and are added to a master list of valid email addresses for future use.


Another technique related to DHAs is to determine whether an SMTP server's MTA has the EXPN (expand) and VRFY (verify) commands enabled. If these commands are enabled, a spammer can use them against you to determine usernames and passwords on your server. Although more complex, this gathering method could be considered a type of directory harvesting attack.

 See SMTP Request for Comment (RFC) 821 for more information about these commands at <http://www.ietf.org/rfc/rfc0821.txt>.

No matter how email addresses are obtained, once they make it to one spam list, it's usually just a matter of time until they get on more and more, making the spam problem exponentially worse.

 There are dozens of email marketing companies on the Internet that sell these email lists to spammers for as little as \$100 for tens of millions of email addresses.

Keep in mind that any time you disclose your email address to anyone, you're risking being added to a spam list. Once your email address is made public, there's basically no way to control how it is used. Of course, we have to give out our email addresses; otherwise, why have email?

 Later in this chapter, I'll show you some methods you can use to give out your addresses while minimizing your chances of getting onto spam lists.

## Identifying Spam

There are a few key characteristics that you and spam filtering applications can use to tell whether an email is spam. Although the following list isn't comprehensive, history shows that if a message contains one or more of these characteristics, it's most likely spam:

- From address and Reply-To address do not match—This characteristic is one of the most common for spam (although there are exceptions, such as when someone uses a work account to send email, but they want recipients to reply to a personal account).
- Sender address and recipient addresses match (as if you've sent a message to yourself)—This sneaky tactic is spammers way of trying to remain anonymous and making the email look legitimate to elementary spam filters because the email addresses appear to be legitimate.
- Recipient addresses that are all very similar (a spam dictionary is used to send these)—When recipient addresses are all similar such as user1@yourdomain.com, user2@yourdomain.com, and so on, the addresses have simply been pulled from a dictionary of all possible email addresses in an attempt to determine your valid address. Sometimes these addresses are harvested via DHAs.
- Large number of spaces in the subject field—This tactic is an elementary attempt to bypass filters that search for specific words.
- Trailing numbers and/or letters in subject field or message body—Random characters are placed in the email to throw off any signature-based spam-detection systems. For example, you might have a subject line of *Summer school Project WOXLWBTKGS*. This method has been extremely popular recently, which is simply a sign that the spammers are in a sort of arms race with spam-filtering solution vendors.
- Misspelled words in the subject line or body—This method is yet another way spammers try to bypass filters that look for specific words in spam messages. An example might be *Fre"e Mortgage Quote*. Although this content is somewhat readable by humans, it is typically undetectable by spam-filtering technologies.
- Sender addresses that are largely numeric (for example, 3347866@some\_spam\_server.net)—These addresses are disposable (I'll talk more about these later in this chapter).
- Celebrity names listed in the subject line—This trick is used to draw people's attention to the email and entice them to open it and learn more about a popular celebrity.
- Predictable phrases (for example, "get rich quick" and "work from home")—These phrases are easily caught by spam filters and the human eye, but still work nonetheless.

The open nature of the Internet and specifically SMTP allows spammers to get by with their shenanigans fairly easily without getting caught. A large portion of spam messages I've seen, especially recently, are unique in some way. Whether the message has a few random characters entered into the subject line or message body or a completely different subject line altogether, spam messages are morphing to change their characteristics. This one small factor is what is keeping traditional anti-spam technologies from being highly effective. These unique spam identifiers are just a few ways that professional spammers try to trick their recipients into opening their emails. They're using a form of social engineering, which I covered in Chapter 2, and exploiting the natural trusting tendency (and curiosity) of human beings to con people into reading their messages.

### Tracking Spam Sources

Have you ever wanted to track where your spam comes from? It's not easy because of the various disguise methods that spammers use, but you might be able to glean enough information from a spam message to track it back to an ISP. The key to doing so is to understand how to read email headers. An email header keeps track of date and time stamps, who sent the email, who is to receive the email, which servers the email has traversed in route to its destination, and so on. The following header is from an actual spam message I received. There was a lot of impertinent information contained in this header that I have removed for the sake of clarity. I have numbered each major line so that we can explore what they mean.

```

1. Received: from lsanca1-ar11-4-60-100-095.lsanca1.dsl-verizon.net
(Not Verified[4.60.100.95]) by mail.principlelogic.com with MailMarshal
(v5,0,3,100) id <B000005f67>; Tue, 10 Jun 2003 00:00:41 -0400

2. Received: from d76.qehl.com [222.173.40.167] by lsanca1-ar11-4-60-
100-095.lsanca1.dsl-verizon.net with SMTP for <kbeaver @
principlelogic.com>; Tue, 10 Jun 2003 00:04:09 -0500

3. Message-ID: <r$5y-b$c0-2uar3-3dg-v4@v74idqfab>

4. From: "Tanisha Hager" <o3k8ehujws@ibm.com>

5. To: kbeaver @ principlelogic.com

6. Subject: Actually reverse a g i n g symptoms! as muj sbrarijju v

7. Date: Tue, 10 Jun 03 00:04:09 GMT

```

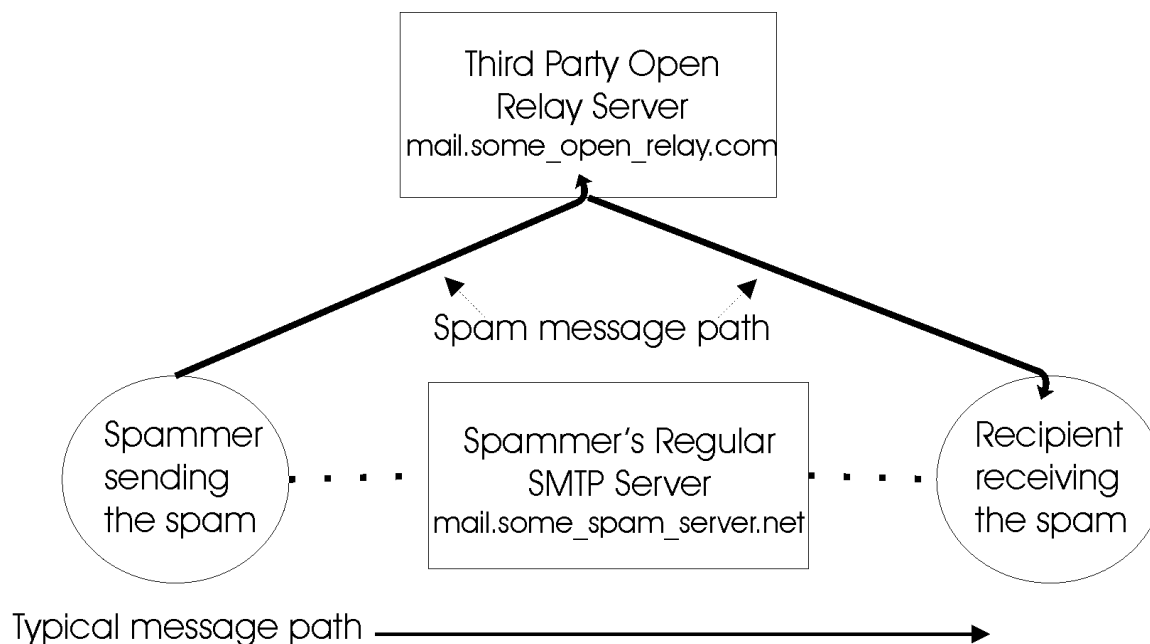
The first line shows the final path that the email took—from the lsanca1-ar11-4-60-100-095.lsanca1.dsl-verizon.net host to my own mail server. The second major line shows the first path that the email took (from the d76.qehl.com host to the lsanca1-ar11-4-60-100-095.lsanca1.dsl-verizon.net host). I put the spaces around the “@” symbol when editing this for my protection. And the two lines that follow were added for each host the email traversed. This information along with the date and time stamps were originally designed to troubleshoot email delivery problems, but they can obviously be used to track spam as well.

👉 Once you have a hostname or IP address of a potential spam source, you can go to <http://www.samspade.org> to discover all the information you *should* need to help track spam sources.

The third major line is a unique message identifier that can be used to track the message from the beginning to the end of its transmission. Next, on line four, is the name and email address of the “sender.” Line five is the email address to which the email was sent. (Again, I put the spaces around the “@” symbol.) Line six shows the email’s subject line. Note the spaces between the letters in the word “aging,” which were added in an attempt to slip past content filters, and the random letters at the end that give it a unique electronic signature that can throw off some signature-detection filters. Finally, line seven is the date that the email was sent.

### ***Sending Through Open Relays***

A lot of spammers use open SMTP relays to send their wares throughout the Internet. They are exploiting what was originally designed into SMTP as a testing and message delivery enhancement mechanism. Email servers complying with the SMTP standard contain a relay feature that allows hosts on the Internet to send email through an SMTP server other than their own. In this situation, the third-party relay server does not contain a local account for the sender and/or receiver of the email but rather more of an innocent and trusting bystander. This setup provides a level of redundancy to help ensure message delivery in case the host’s main SMTP server is down. Figure 3.2 shows the path an email takes when traveling through an open relay.



**Figure 3.2:** Email message path through an open relay.

Your email server should have an easily accessible option to turn off SMTP relay. You can use the following steps to test your open relay before you reconfigure your server. Keep in mind that your SMTP server may require additional steps. If so, you can refer to your email software administrator's guide, search the Internet for the specific commands, or better yet, use an automated tool as I describe later:

1. Telnet to your email server on port 25 either from a command prompt by typing

```
telnet your_mail_server_name_or_address 25
```

or using your favorite telnet application within your desktop environment. You may need to enable local echo of characters in your telnet application in order to see what you're typing.

2. You should see some sort of welcome banner. Type

```
mail from:yourname@yourdomain.com
```

and press Enter. You should receive a "250 OK" or similar message.



Technically, you need to use a HELO *yourdomain.com* or EHLO command to initiate the session with your server before entering the mail from: command. Please refer to your SMTP server's documentation for the specific commands.

3. Type

```
rcpt to:yourname@yourdomain.com
```

and press Enter. You should receive a "250 OK" or similar message.

4. Type

```
data
```

and press Enter. You should receive a "354 enter mail" or similar message.

5. Type

```
Subject:This is a relay test
```

and press Enter.

6. Type a couple of words for the body of the message, and press Enter. Type a period on a separate line by itself to end the message. You should receive a "250 OK" or similar message.

7. Check your email account to see whether the actual email was relayed.

Figure 3.3 shows these steps in action.



```

telnet - HyperTerminal
File Edit View Call Transfer Help

220-mail.principlelogic.com ESMTP Merak 5.3.2; Thu, 12 Jun 2003 14:55:10 -0400
220 WARNING! - This is a private system. All use is monitored and recorded. Any
unauthorized or malicious use may result in criminal and/or civil prosecution.

mail from:spammer@somespamserver.net
250 2.1.0 <spammer@somespamserver.net>... Sender ok

rcpt to:anyone@principlelogic.com
250 2.1.5 <anyone@principlelogic.com>... Recipient ok

data
354 Please start mail input.

Subject:This is a quick a dirty relay test
I think it worked!

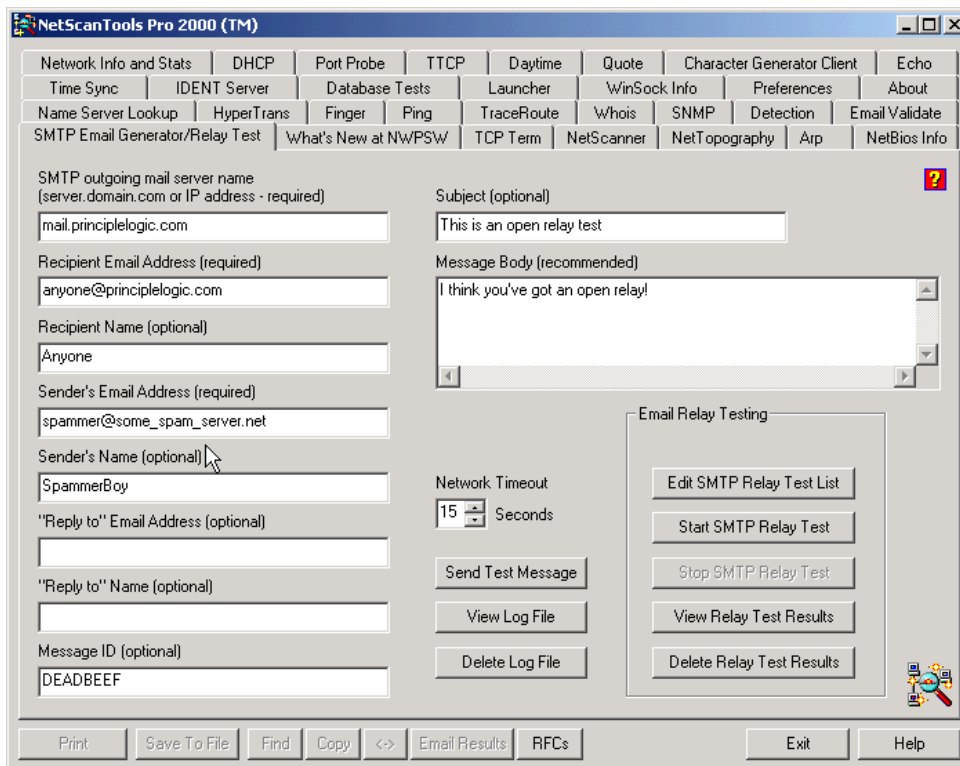
.
250 Mail queued for delivery.

Connected 0:01:44 ANSIW TCP/IP SCROLL CAPS NUM Capture Print echo

```

**Figure 3.3: A quick and dirty SMTP relay test.**

If you'd like to test an open relay, there are several Web sites that allow you to do so (such as <http://www.abuse.net/relay.html>). In addition, there are freeware and commercial tools that have this functionality, such as ExtraLAN's NetScan Tools Pro (see Figure 3.4). These tools are typically easier and leave less room for error than the manual process I just described.



**Figure 3.4: Open relay test using an automated tool.**



### **Additional Spam Propagation Techniques**

There are various other tricks that spammers use to send spam messages and slip by spam filters such as:

- Changing the date and/or time on their email software to make spam appear at the front of the email list. I've even seen spam that comes in with a really old date and ends up showing up way at the bottom of an Inbox. These tricks may not be a problem for Inboxes with only a few messages, but for those of us who have dozens or more messages, it works pretty well.
- Using variants of domain names so that the spam skips around and bypasses filters. For example, mail.some\_spam\_server.net might be changed over to mail.sumspamsrver.net or mail.some\_spam\_server1.net.
- Using disposable email addresses such as %67113#@some\_spam\_server.net. Any attempts at filtering these types of addresses is typically futile.
- Employing email addresses in domains that most people don't want to (or cannot) block (such as aol.com, hotmail.com, yahoo.com, and so on).
- Making use of the same domain name and changing the server's IP address and DNS address record.
- Typing with very large fonts to try to catch the recipient's eye.
- Using graphics instead of text to display the message in the email body.
- Forging email headers to make it look as if the spam came from someone else.
- Inserting HTML comments between letters and words that people won't see but can trick spam filters.
- Using offshore servers that are not yet blacklisted to send from or relay through.

### **Spam Elimination Techniques**


Before you can effectively block spam, you need to determine who is going to be responsible for developing and managing the appropriate anti-spam policies, procedures, and technologies.

Unfortunately, like malware, spam is everyone's issue to contend with. Luckily, with today's network and server-based anti-spam technologies, you can catch most spam before it ever gets to the desktop level. However, no matter how good your spam filters are, a certain amount of spam will still get through. Thus, end users need to be made aware of how to handle and report spam when they receive it.

The responsibility for spam prevention and management should lie with IT and/or information security. Sure there are a lot of generic IT issues related to spam, but like I mentioned earlier, spam can certainly create a lot of security issues as well. If you have to choose, choose both IT and security (assuming they're separate), and make the roles and responsibilities as far-reaching as possible with the head of IT or security ultimately responsible.

### **Disable Open Email Relays**

Got an open relay? If so, then disable it. There should be no reason that an email server needs to relay public emails. The only logical time relay should ever be enabled is when internal hosts or other mail servers need to send out through a primary email server. Enabled public relay not only lets spam through, but it is also an invitation to become listed on blacklists of known open relays (I'll talk more about blacklists in a moment). Most email servers let you turn on relay for internal hosts without a problem. Keep in mind that many email servers come with SMTP-relay enabled, although many vendors are now configuring it as disabled by default.

 For instructions about how to disable SMTP relay on various email server platforms, check out <http://www.mailabuse.org/tsi/ar-fix.html>.

### **Filtering Methods**

A just-say-no approach to spam doesn't work. Simply asking your users to just delete their spam messages as they come in is an expensive exercise in futility. It certainly won't eliminate the spam problem long term. Policies against spam aren't enough either. They can help discourage and prevent spam coming from internal users, but external spammers are certainly not going to care about them! Your best bet is to avoid spam whenever possible by using anti-spam technologies.

There are various spam-filtering methods available, but there is often a tradeoff between efficiency and accuracy. Some solutions are better than others for detection and preventing false positives. For example, scanning messages for keywords can lead to a high number of false positives and blocked emails that are otherwise legitimate. In today's business environment, even one false positive is unacceptable. Filtering methods must continuously improve to fight off the new spamming techniques constantly being developed. As I mentioned earlier, this one-upmanship has led to an arms race between spammers and anti-spam experts and vendors.

The following sections describe various spam-filtering methods from a high-level perspective as well as the pros and cons of each. If your spam-filtering applications support some or all of these methods, it might behoove you to enable them depending on your specific needs and the amount of risk you're willing to take with false positives. As with any information security infrastructure, a defensive in-depth strategy works best. In other words, you cannot simply rely on one spam defense mechanism and expect to achieve maximum results. You must implement various layers of protection that work in unison to provide the greatest possible protection.

## Blacklist Filtering

Blacklists are lists of IP addresses and/or domain names of known spammers and open SMTP relays. You can have your own blacklist or utilize a third-party blacklist. There is no centralized blacklist organization, but several organizations have taken it upon themselves to become central repositories of blacklist databases. These third-party blacklists are typically supported by volunteers and monetary donations. The pros of blacklist filtering include:

- Integrate easily with most email servers and spam-filtering applications
- Easy to set up and use—sometimes as simple as a check box on your email server or spam-filtering application
- Free or low-cost solution
- Option to use a third-party database that is maintained by someone else

The cons of blacklist filtering include:

- May block legitimate messages (although many people think it's well worth losing a few emails here and there if it means that their IT and user resources are free from most spam)
- Filtering rates may not be as high as other filtering methods
- Can be easy to circumvent as spammers constantly change their domains and IP addresses
- Must rely on other sources to update the third-party blacklist databases
- A third-party blacklist's definition of spammers may not be the same as yours
- Blacklists come and go—some of the less-popular lists might not be around long-term and are therefore not a dependable solution
- Some blacklists can be very difficult to get unlisted from—perhaps good for spammers but bad for legitimate users

The following list offers a few real-world examples of blacklist services:

- Mail Abuse Prevention Systems (MAPS) at <http://www.mailabuse.org>
- Open Relay Database (ORDB) at <http://www.ordb.org>
- Distributed Server Boycott List at <http://www.dsbl.org>
- Spam Prevention Early Warning System at <http://www.spews.org>
- Spamhaus at <http://www.spamhaus.org>
- SpamCop at <http://www.spamcop.net>

When signing up for a blacklist service, be sure to find out the organization's criteria and method for adding and removing domains and IP addresses as well as how quickly they do it. Another issue to consider when signing up with a new ISP or for new hosting or co-location services, check the IP address block that you are going to be assigned (and those IP addresses that are close in range to yours) against your favorite blacklist database to make sure that you're not already blacklisted before you even connect to the Internet.

☞ If you get accidentally blacklisted for whatever reason, follow these steps to ensure that you're cleared as soon as possible:

- First, notify your critical contacts such as customers and business partners.
- Have your critical contacts send messages to a backup Web-based or alternative domain address.
- Find out which blacklist(s) your servers are listed on and follow their instructions *verbatim* to get removed. One of my favorite sites to check for blacklisting is the Osirusoft database located at <http://relays.osirusoft.com/cgi-bin/rbcheck.cgi>. This database searches most of the well-known blacklist databases, making your search much simpler and quicker.
- Monitor the blacklist to ensure that you've been removed to prevent future blockages.

## Whitelist Filtering

Whitelists are simply lists of email addresses, IP addresses, or domain names that you trust to send messages into your system. This list is typically a self-maintained database of organizations and people you regularly communicate with. Real-world examples of whitelist filtering include major ISPs and Web-based email providers (Yahoo and AOL use whitelists). Most email client and server programs allow you to create whitelists. The pros of whitelist filtering include:

- Even if you have a blacklist that has blocked an entire domain, you can permit certain subdomains to send email to you using a whitelist.
- Some whitelist technologies built-in to certain anti-spam products can “learn” who should be on the whitelist. These methods monitor the recipients of outbound emails and add those recipients after a certain number of outbound messages have been sent to the same address or domain.
- Guarantee delivery of legitimate emails
- You have control over the database
- Very high filtering percentage—virtually 100 percent
- Tight integration with most email servers and spam-filtering applications

The cons of whitelist filtering include:

- There's a good chance that legitimate email is not going to make it into your system because you may have forgotten to allow certain addresses or domains.
- Can take a lot of time and resources for initial creation and ongoing maintenance of the whitelist.

## Signature Filtering

Signature filtering works by performing a checksum or other cryptographic hash (such as MD5) on individual emails, filtering out the ones that match a known spam fingerprint. The pros of signature filtering include:

- Highly accurate because there is theoretically no way two messages can have the same electronic signature.
- Is built-in to more and more spam-filtering applications.

The cons of signature filtering include:

- Processor intensive, which can slow email delivery or possibly create an email DoS situation on high-volume networks.
- Can be easily circumvented by inserting random characters into spam messages because doing so would completely change a message's electronic signature.
- There are so many spam options and so many ways to alter spam messages, a signature-based system can have trouble keeping up. However, vendors have come up with ways to fight back against this type of circumvention, including determining whether a certain percentage of the signature matches or common words or other characteristics match.

Real-world examples of commercial services and products that offer signature filtering include Singlefin's spam filtering service (<http://www.singlefin.net>) and CipherTrust's IronMail (<http://www.ciphertrust.com>). Open source services and products include Vipul's Razor (<http://razor.sourceforge.net>), Distributed Checksum Clearinghouse (<http://www.rhyolite.com/anti-spam/dcc>) and Cloudmark's Spamnet (<http://www.cloudmark.com>).

## Heuristics Filtering

Heuristics filtering works through trial and error checking for various known spam characteristics such as well-known phrases, forged email headers, and larger than normal fonts, which may signal that a message is spam. Heuristics filtering provides a high level of accuracy and many newer versions of anti-spam products are incorporating this type of filtering. However, this filtering method can be more processor intensive than other methods and it requires more maintenance as detection techniques can fall behind the development of new spam methods.

## Email Header Filtering

In a nutshell, email header filtering looks for malformed information in each email's header section. This method sometimes includes reverse DNS lookups to verify that the sending domain or address is legitimate. The benefits of email header filtering are that anti-spam products support this type of filtering and it is fast and efficient because the entire message body is not analyzed. The drawbacks to email header filtering include:

- May block legitimate emails coming from legitimate servers that are misconfigured or have no reverse DNS lookup available.
- Spam servers or open relays that are properly configured can sneak past this form of filtering.
- Performing reverse DNS lookups for every inbound email can create a massive amount of traffic that can slow an Internet connection and eat up resources on DNS servers.

## Content Filtering

Content filtering applications look at specific words and attachments to determine whether a message is spam. Techniques include scanning for vulgar language, images based on human skin tones, popular spam phrases, and image tags pointing to URLs that are known to be bad. Every content-filtering application offers different options for the email administrator. Figure 3.5 shows some of the content filtering options in Singlefin's email protection service. Notice that this solution lets you customize filters to search for certain character sets as well as words and phrases in the message subject and body.

Organization:  Domain:

**This is the default domain. Changes will affect all domains using the organization default**

<b>Content Filtering:</b>	<input checked="" type="checkbox"/> <b>Enable</b>
Character Sets:	Character Set Block: <input type="text" value="afrikaans [af]"/> <input type="button" value="add"/> Character Sets Blocked:
Subject Block:	<input type="text"/> <input type="checkbox"/> Exact Match <input type="checkbox"/> Case Sensitive ex: free mortgage loan <input type="button" value="add subject block"/>
Subjects <b>blocked</b> :	Free Money EM <a href="#">edit / delete</a> <a href="#">edit / delete</a>
Message Body Text Block:	<input type="text"/> <input type="checkbox"/> Exact Match <input type="checkbox"/> Case Sensitive ex: this is not spam <input type="button" value="add message body text block"/>
Message Body Text <b>blocked</b> :	this is not spam EM <a href="#">edit / delete</a> <a href="#">edit / delete</a>


Figure 3.5: Singlefin's content-filtering options.

The benefits of content filtering include:

- Offers a high level of customization by the end user or network administrator.
- Available in practically all email client and server software and anti-spam applications.


The cons on this filtering method include:

- Can generate a large number of false positives.
- There are many well-known ways to circumvent content filtering.
- Many content-filtering solutions focus on antivirus and Web surfing control rather than spam prevention
- Large ISPs and Web-based email providers are performing content filtering that result in an increasing amount of legitimate email that never makes it to our Inboxes.
- Email clients use the “flag as spam” feature to block legitimate messages for which they previously opted-in.

 According to Assurance Systems, one in six opt-in permission emails is considered spam and blocked. Check out the company's statistics at <http://www.assurancesys.com>.

## Bayesian Filtering

Spammers have found a way around every spam filtering technique—until now. Bayesian filtering analyzes the language used in emails and assigns probabilities to certain spam-related words and phrases and certain non-spam related words and phrases referred to as tokens. This filtering method fine tunes the probabilities and its own filters over time based on an individual email address or user.

 As defined by the Merriam-Webster online dictionary, Bayesian means “being, relating to, or concerned with a theory (as of decision making or statistical inference) involving the application of Bayes' theorem and the use of probabilities based on prior knowledge and accumulated experience.”

The pros of Bayesian filtering include:

- Extremely high level of accuracy
- Fewer false positives
- Looks for the good and the bad in emails, not just the bad
- Virtually maintenance free
- All of the spamming tricks, even new and undocumented ones, are used to determine whether a message is considered spam

The cons of this filtering method include:

- Doesn't scale well to the network/perimeter level; needs to be used at the desktop level so that the individual email user's email behavior can be correctly analyzed
- Requires a larger amount of processing power compared with other filtering methods

You can find open source Bayesian filtering products from SpamAssassin (<http://www.spamassassin.org>) and SpamBayes (<http://spambayes.sourceforge.net>).

### **Additional Ways to Fight Spam**

Besides disabling open SMTP relays and employing various filtering techniques, there are a few other methods you can implement at the client and server level and outside your network to block unwanted spam:

- Don't reply to spam messages that con you into believing that if you reply and ask to be removed from the spam list, you'll actually be removed. This method of fighting spam rarely works. Instead, replying to spam simply confirms that your email address is valid and the spammers will use your address even more in the future.
- Use a free Web-based account when you're posting public messages or giving out your email address to public sources. These are typically filtered and won't clog up your regular Inbox.
- Enable any spam controls made available for free by your ISP.
- Use a third-party filtered email account. You can sign up for one of these free accounts and use it instead of your real address when giving out your email address. All email will then be sent to this address and be filtered before being forwarded on to your actual account. There are several options available on the Internet. One of the more popular services can be found at <http://www.despammed.com>.
- Use a unique email account on your domain for every site you sign up with, for every month of the year, and so on. Be aware that this method requires you to alert everyone who has your address that you have a new address, which can add up to quite a bit of maintenance—you don't realize how many people you've given your address to until you start trying to contact everyone to give them the new one!
- Use a disposable email address from an external provider; the following list provides some of the disposable email address providers:
  - Emailias.com at <http://www.emailias.com/>
  - MailShell at <http://www.mailshell.com>
  - Sneakemail at <http://www.sneakemail.com/>
  - SpamCon Foundation at <http://dea.spamcon.org>
  - Spamex at <http://www.spamex.com/>
  - SpamGourmet.com at <http://www.spamgourmet.com/>
  - SpamMotel at <http://www.spammotel.com/>
  - TrashMail at <http://trashmail.net>





- A method that won't necessarily block emails but can provide a way to track the origin of spam messages is to use *plussed* email addresses when registering with Web sites that you're not familiar with (or otherwise giving out your email address). For example, you can enter the following: `yourname+untrusted_site.com@yourdomain.com`. Your email server will support these types of email addresses, and this method can be an effective way to monitor spam.
- Capture spam for further analysis (and addition to blacklists) by creating dummy email accounts and post them on Web pages or other places where you know they will be captured by spam-harvesting software. You could also set up servers as spam traps by enabling SMTP relay but instead of actually relaying the messages, you could forward them to an internal account for further analysis.
- Use a community or P2P-based spam judging solution such as Cloudmark's Spamnet. This system let's the P2P members "vote" on whether a message is spam. Once enough P2P members vote that a message is indeed spam, the Spamnet system will block those messages in the future.
- Complain to the postmaster ([postmaster@some\\_spam\\_server.net](mailto:postmaster@some_spam_server.net)) at the spam origin.
- Consider using an ISP that has an acceptable usage policy against spamming so that you know the ISP will (hopefully) take measures to fight spam.
- If you have to post email addresses on a Web site or in your email signature line, consider using the following methods:
  - Use special HTML character coding in place of the ":", "@", and "." in mailto links and HTML-based email signatures. For example, for a link containing `mailto:yourname@yourdomain.com`, you would enter the following
 

```
mailto&#58;yourname&#64;yourdomain&#46;com
```

However, this solution isn't foolproof, as some spambot programs can detect this type of obscurity.
  - As an alternative to the previous item, you can enter an email address in HTML format as follows:
 

```
<table>
<tr>
<td>yourname</td>
<td>@</td>
<td>yourdomain.com</td>
</tr>
</table>
```
  - Put spaces before and after the "@" character or use a different set of characters (such as <a>) in your reply-to address or in your email signature. Humans will know what the difference is, but doing so can throw off spam-harvesting software.

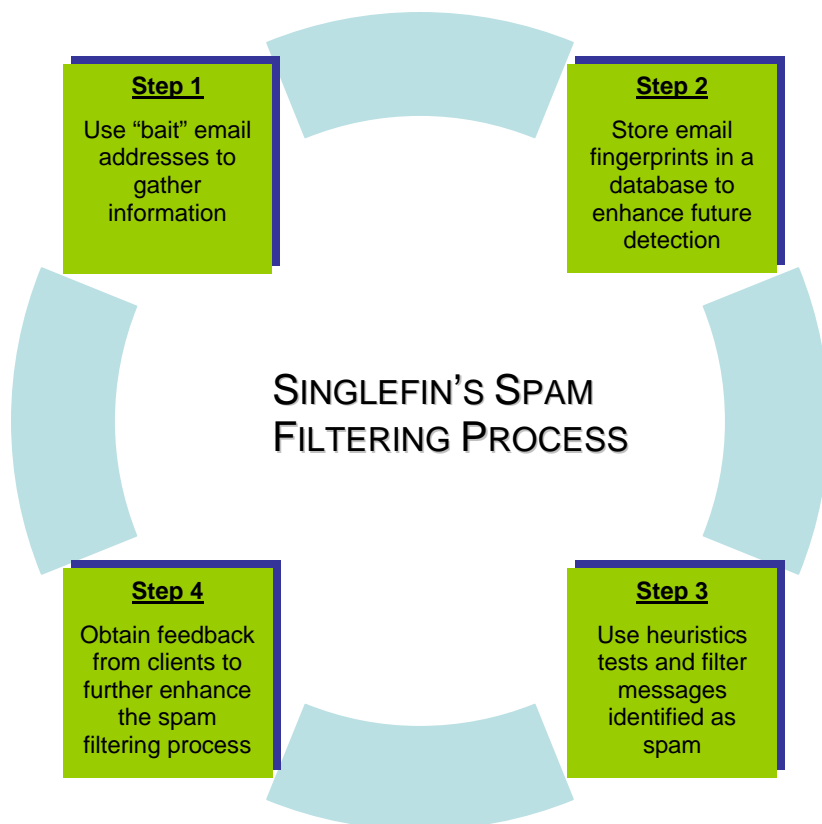
- Append @nospam or .nospam to your address and put a signature line at the bottom of your message instructing the user to remove the obvious addition to your address before replying.
- Don't use an out-of-office auto reply on your email client, or if possible, configure it for use only with important clients or people in your address book.
- Keep your end users educated about what to do and if spam slips past your filters
- Keep yourself educated about the latest spam-blocking techniques
- Create a policy and set of procedures for forwarding spam to outside entities that can track and notify spammers' ISPs or even the Federal Trade Commission, which can monitor for fraud.

### ***Layered Anti-Spam Defenses***

Spam filtering can take place at the client level, the server level, the network perimeter level, the third-party ASP level, and even at the ISP level. As with any other information defense mechanism, a true multi-layered process works best for higher accuracy when trying to block spam. Of all the filtering methods we've explored, there's not one single best defense that you should deploy. Rather, you should consider using several if not most of them—especially the filtering methods—at the level you choose for your defenses (server, perimeter, ASP, and so on).

Be careful when installing and fine-tuning your spam-filtering application. You might get more filtering capabilities than you bargained for. For example, in mid-2003, a “technical change” of the AOL spam filters caused millions of emails destined to AOL subscribers to be blocked from several large ISPs in the United States supposedly for as many as 10 days. Key points to keep in mind when configuring your spam-filtering application:

- Managing spam on each user's computer can be an exercise in futility especially for larger organizations. Spam filtering is not usually enabled at the client level by default, but can be enabled in the email client fairly easily or by installing anti-spam software. Consider that IT staff and end users might end up spending more time than necessary “tweaking” their spam filters.
- ASP spam solutions can be very beneficial for smaller organizations that have little or no IT staff and for larger organizations that simply cannot justify dedicating resources to managing spam. Just make sure that you do your homework in selecting the right vendor. There are business-continuity issues that must be taken into consideration (such as what you would do if the provider went out of business suddenly and your email stopped flowing). Another benefit of ASP solutions is that they don't allow email to get to your local systems and eat up bandwidth, storage space, and so on. Figure 3.6 displays the layered spam elimination steps used by Singlefin.



**Figure 3.6:** A high-level view of the spam filtering process by Singlefin.

- There are various dedicated spam appliance solutions that work outside of your email server. These are a good idea if you have a high volume of email, as they can help take a major processing load off your email servers freeing up their resources to do what you installed them to do—send and receive legitimate emails.
- There are various server-based spam solutions that either run as a separate application or tightly integrate with your email server. Keep in mind that you'll need extra server horsepower (processor, memory, disk storage, and so on) to handle both your email server and your anti-spam applications. If possible, you can eliminate this need by installing your anti-spam product on a separate server. Also, keep in mind that running your email server and anti-spam application on the same server can introduce a single point of failure.

### **What to Look for in Anti-Spam Solutions**

There are dozens of anti-spam solutions on the market today. Some are more feature-rich than others. The following list provides some characteristics that you should look for when searching for a good solution for your email system:

- Filters at least 90 to 95 percent of spam
- Requires minimal administration
- Keeps spam from reaching your email server (if not your entire network)
- Uses various methods of spam filtration
- Provides the ability to use only certain filtering methods based on your specific needs
- Has minimal impact on email delivery time
- Has minimal computing requirements for host-based spam solutions
- Integrates monitoring, reporting, and archiving
- Allows for customization of filtering rules including adding your own, turning specific ones off, and so on
- Employs customizable message handling (such as deleting, forwarding, and labeling the subject line and signature/footer areas of emails)

#### **Anti-Spam Tools and Resources**

The following list provides popular server, network, and ASP-based anti-spam solutions. However, this list is not complete and doesn't offer any of the available desktop-based spam filtering products.

##### **Vendors and Tool Providers**

8e6 Technologies at <http://www.8e6technologies.com/>

ActiveState at <http://www.activestate.com/>

Activis at <http://www.activis.com/en/>

Agilera at <http://www.agilera.com/>

AppRiver at <http://www.appriver.com/>

Beginfinite at <http://www.beginfinite.com/>

BlueCat Networks at <http://www.bluecatnetworks.com/>

Brightmail at <http://www.brightmail.com/>

BVRP at <http://www.bvrpusa.com/>

Caledonia at <http://www.caledonia.net/>

CipherTrust at <http://www.ciphertrust.com/>

Clearswift at <http://www.clearswift.com/>

Elron Software at <http://www.elronsoftware.com/>

ePrivacy Group at <http://www.eprivacygroup.com/>

Escom at <http://www.escom.com/>

Gordano at <http://www.gordano.com/>

IntelliReach at <http://www.intellireach.com/>  
LAN-ACES at <http://www.lanaces.com/>  
MailFrontier at <http://www.mailfrontier.com>  
MailWise at <http://www.mailwise.com/>  
Mail-Filters.com at <http://www.mail-filters.com/>  
MessageLabs at <http://www.messagelabs.com/home/default.asp>  
Mirapoint at <http://www.mirapoint.com/>  
NetIQ at <http://www.netiq.com/>  
Network Associates at <http://www.nai.com/us/index.asp>  
OpenHand at <http://www.openhandhome.com/>  
Singlefin at <http://www.singlefin.net/>  
SurfControl at <http://www.surfcontrol.com/>  
Sybari Software at <http://www.sybari.com/home/>  
Syntegra at <http://www.us.syntegra.com/>  
The Messaging Architects at <http://www.gwtools.com/>  
Vircom at <http://www.vircom.com/Enterprise/>

#### Resources

Coalition Against Unsolicited Commercial Email at <http://www.cauce.org/>  
Coalition Against Unsolicited Bulk Email, Australia at <http://www.caube.org.au/>  
Death to Spam at <http://www.mindworkshop.com/alchemy/nospam.html>  
Elsop's Anti-Spam Page at <http://www.elsop.com/wrc/nospam.htm>  
FAQ about tracing an email message at <http://www.cs.uu.nl/wais/html/na-dir/net-abuse-faq/spam-faq.html>  
IETF Anti-Spam Research Group at <https://www1.ietf.org/mailman/listinfo/asrg>  
Junkbusters at <http://www.junkbusters.com/>  
Novell anti-spam tips at  
[http://www.novell.com/coolsolutions/gwmag/features/tips/t\\_tips\\_for\\_spam\\_gw.html](http://www.novell.com/coolsolutions/gwmag/features/tips/t_tips_for_spam_gw.html)  
Reading email headers at <http://www.stopspam.org/email/headers/headers.html>  
SpamArchive at <http://www.spamarchive.org/>  
SpamCon Foundation at <http://www.spamcon.org/>  
SpamLaws at <http://www.spamlaws.com/>  
The Web Robots Database at <http://www.robotstxt.org/wc/active.html>

## Summary

In this chapter, I've covered why spammers spam, how they get your email addresses, tricks they use to sneak their junk into everyone's mailboxes, as well as some proven technologies and techniques to help with spam reduction. The spam problem we currently face requires more than just instructing end users to simply hit Delete and get on with their day. It is becoming a more than a nuisance and affects everyone's productivity, ultimately placing an undue financial burden on organizations.

An oft-overlooked potential positive side-effect of spam elimination is improved morale among employees. They will not only feel more positive about their productivity but will know that their organization is taking serious steps to eliminate emails that can be outright offensive. The ultimate solution may eventually require that we simply deny all email and rely on whitelists similar to the buddy lists we use in our instant messaging software. Let's hope it doesn't get to that point. I truly believe that with a certain amount of time and effort installing and maintaining the proper anti-spam technologies and educating yourself and your users about existing and evolving blocking techniques, we can bring spam to its knees.

In Chapter 4, we'll explore content filtering in more depth. I'll go into detail about how email content filtering systems work, the issues surrounding employee monitoring and employee privacy, and various security policy considerations related to content filtering.