



realtimepublishers.com™

The Definitive Guide™ To

Controlling Malware, Spyware, Phishing, and Spam

McAfee®
Proven Security™

Dan Sullivan

Chapter 6 Spam in the Enterprise	107
Email Operations and Spam Techniques	107
Early Spam and Reactions to It.....	107
The Basics Steps in Email Operations	108
Email Clients.....	108
Email Servers	109
Vulnerabilities in Email Infrastructure	110
Economics of Spamming	112
Costs and Revenues of Spamming.....	113
Negative Externalities of Spam	114
Costs Incurred by Others	114
Distorting the Supply/Demand Balance.....	115
Correcting for Negative Externalities: Government Regulation	116
Beneficiaries of Spam	119
Spam Management.....	120
Detection and Determination	121
Integrity Analysis.....	121
Heuristic Detection	121
Content Filtering	122
Blacklists and Whitelists.....	122
Self-Tuning	123
Bayesian Filtering	123
DNS Block Lists	124
Actions in Response to Spam	125
Blacklist and Whitelist Management	125
Quarantine Access, Search, and Management.....	126
Evaluating Anti-Spam Systems	126
Catch Rates	127
False Positives.....	127
Manageability and Reporting.....	127
Summary	128

Copyright Statement

© 2006 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.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 Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.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.

Realtimepublishers.com and the Realtimepublishers 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 Realtimepublishers.com, please contact us via e-mail at info@realtimepublishers.com.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtimepublishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 6 Spam in the Enterprise

Spam, or unwanted and unsolicited email, in the enterprise unnecessarily taxes IT resources. Unlike its kin, phishing scams, spam itself is not a direct threat to security; rather the damage it causes is the result of the fact that it consumes network bandwidth and storage as well as wastes employees' time. As part of broader compliance initiatives, companies may be required to archive all email messages for extended periods of time, so even if spam is deleted by end users, it could continue to consume storage for years to come.

 For more information about phishing, see Chapter 5.

This chapter begins by examining the basic operations of mass emailing and discussing how spammers exploit weaknesses in email protocols. Next, it addresses the economics of spam and the attempts to control spam through legislation. Although helpful, legislation has not stopped spam and likely will not. Technology is therefore crucial to managing spam. This chapter includes a review of spam management techniques and concludes with some guidelines for evaluating anti-spam systems.

Email Operations and Spam Techniques

When unsolicited mass mailings and newsgroup postings began, spammers sent messages the way any other user would send messages. There was no attempt to hide its source or pretend to be something other than an advertisement. By the early 1990s, spammers began using programs to automatically post messages to multiple Usenet newsgroups. Today, legislation, such as the CAN-SPAM Act in the United States, attempts to curb spam and has prompted spammers to use techniques that hide their identities and the origins of their messages. These techniques depend on exploiting a combination of vulnerabilities in email protocols and insufficiently protected computers on the Internet. The following discussion examines early examples of spam to demonstrate that this phenomenon is not new.

Early Spam and Reactions to It

Spam is so prevalent today it is difficult to imagine an email user becoming upset about receiving an unsolicited message. During the early years of the Internet, and while its predecessor the Arpanet was in use, small groups of users generally adhered to a common understanding of the proper use of distributed applications such as email, listservs, and newsgroups. (And, in the case of the Arpanet, which was run by the United States military, there were formal rules governing its use as well.) When a user violated this understanding, there was usually a substantial negative response by others in the user community.

For example, in 1978, a Digital Equipment Corporation (DEC) sales representative sent an email to a large number of Arpanet email addresses encouraging recipients to attend a presentation on a new line of computers from DEC. This act was viewed as a flagrant violation of Arpanet regulations as well as common practice. It generated a strong anti-advertising response from other users.

 For the original message and some of the reactions to this early spam, see Brad Templeton's "Reaction to the DEC Spam of 1978" <http://www.templetons.com/brad/spamreact.html>.

Another early spam episode stemmed from a posting from two lawyers at the Canter & Siegel law firm to about 6000 newsgroups. (Newsgroups allow users to subscribe to and receive messages about a particular topic.) In the posting, the attorneys offered their services to anyone who wished to participate in a United States government lottery of "green card" work permits for foreign nationals. Newsgroup users shared a common understanding that newsgroup postings should be relevant—even relevant commercial announcements were allowed. This posting would have been appropriate on newsgroups about work permits, visas, and labor law, but appearing in completely unrelated newsgroups created a storm of protest.

Canter and Siegel received 20,000 inflammatory emails as well as numerous junk faxes. At least two Internet service providers (ISPs) terminated their service because of the volume of network traffic they were generating. Still, they, like today's spammers, were undeterred. The *New Scientist* magazine quotes Canter and Siegel's spammer's manifesto dismissing a common standard for using Internet resources:

The only laws and rules with which you should concern yourself are those passed by the country, state, and city in which you truly live. The only ethics you should adopt as you pursue wealth on the (information superhighway) are those dictated by the religious faith you have chosen to follow and your own good conscience (Source: Charles Arthur, "A Spammer in the Networks," *New Scientist*, November 1994 as found at <http://www.kkc.net/cs/new-sci1.txt>).

The spammer's credo as argued by Canter and Siegel directly contradicted the attitude of many early Internet users. As legitimate business use of the Internet has evolved, we have witnessed the rise of legitimate online advertising and a general acceptance of it. Spam, however, is still seen as largely an inappropriate use of Internet resources.

To realize the benefits of mass emailing without bearing the consequences, spammers have turned to exploiting vulnerabilities in the email system. Before discussing those vulnerabilities in detail, a brief discussion of email operations is in order.

The Basics Steps in Email Operations

Internet email systems are composed of two types of programs: clients and servers, sometimes called user agents (UA) and message transfer agents (MTAs), respectively (see Figure 6.1).

Email Clients

Client programs allow users to create, send, and manage email messages; Microsoft Outlook, Eudora, and Netscape Mail are examples of email clients. They typically provide features to organize messages into folders, offer track lists of email addresses, enable rules for categorizing messages, and provide related functions. Another core function is to coordinate the sending and receiving of messages from email servers.

Email Servers

A mail server is a transfer agent—it moves messages from the sender to the receiver's email client. Mail servers run the Simple Mail Transfer Protocol (SMTP), which listens for messages on TCP port 25. When a message arrives, the mail server examines the header information to determine the recipient, for example someuser@abc.com. In this example, the email should be sent to the mail server at domain abc.com. To do so requires mapping from the domain name to an IP address, so a domain name services (DNS) server is queried.



Figure 6.1: Mail clients and servers work in conjunction to deliver mail from senders to receivers.

DNS servers maintain several mapping records. The ones associated with email servers, MX records, identify the IP addresses of the server that should receive email messages. As Figure 6.1 shows, once the sender's email server has the IP address from the Mail eXchanger (MX) record, the message can be sent on to the recipient's email server.


In some cases, the message passes through multiple email servers. For example, when a user's account is set to forward messages, the server will send the message to the mail server in the domain specified by the forwarding address. Another case when multiple mail servers are used is when mail servers are configured as *open relays*. In this configuration, the server accepts and transfers messages on behalf of any user. Open relay stems from methods of transfer used when constant, high-bandwidth Internet connections were not commonly available. Today, open relay servers are abused by spammers and are therefore not recommended. Open relays are just one of a number of vulnerabilities found in today's email infrastructure.

Vulnerabilities in Email Infrastructure

Vulnerabilities in the email system occur in both client applications and with email protocols. In the case of email clients, vulnerabilities emerge as additional features are added to make client applications more functional and integrated with other desktop applications. The vulnerabilities in SMTP are the result of the simplicity of this protocol, which has been widely used for decades.

Email Client Vulnerabilities

SMTP and early email clients were designed for sending and receiving only text-based emails. Email clients have become more feature-rich as well, supporting scripting languages, address books, and integration with other desktop applications. Although certainly useful, these additional functions have also introduced vulnerabilities into mail clients that have been exploited by viruses, worms, and other forms of malware.

 For more information about worms, viruses, and malware in general, see Chapter 3.

One technique used by spammers exploits scripting languages supported by email clients to collect email addresses from address books. Spammers (and malware developers) can use this technique to either harvest addresses or propagate messages from the victim's account. Other vulnerabilities in email systems result from design choices in email protocols.

Protocol Vulnerabilities

A number of vulnerabilities in SMTP stem from the fact that it trusts participants in message exchanges. The open relay configuration is an example. The purpose was to provide a transfer service for others when it was not always practical to send messages directly between any two email servers. The underlying assumption was that anyone sending messages through an open relay had a legitimate purpose and was not abusing the service to the point that it hampered others.

In the case of client-to-server transmissions, a client can send a message purportedly from a user, which, in fact, is sent by someone else. A service extension to SMTP, SMTP-AUTH, has been developed that provides a mechanism to force users to authenticate before sending messages through the server. Many email servers and clients now support this extension; Figure 6.2 shows a typical type of client configuration dialog box.

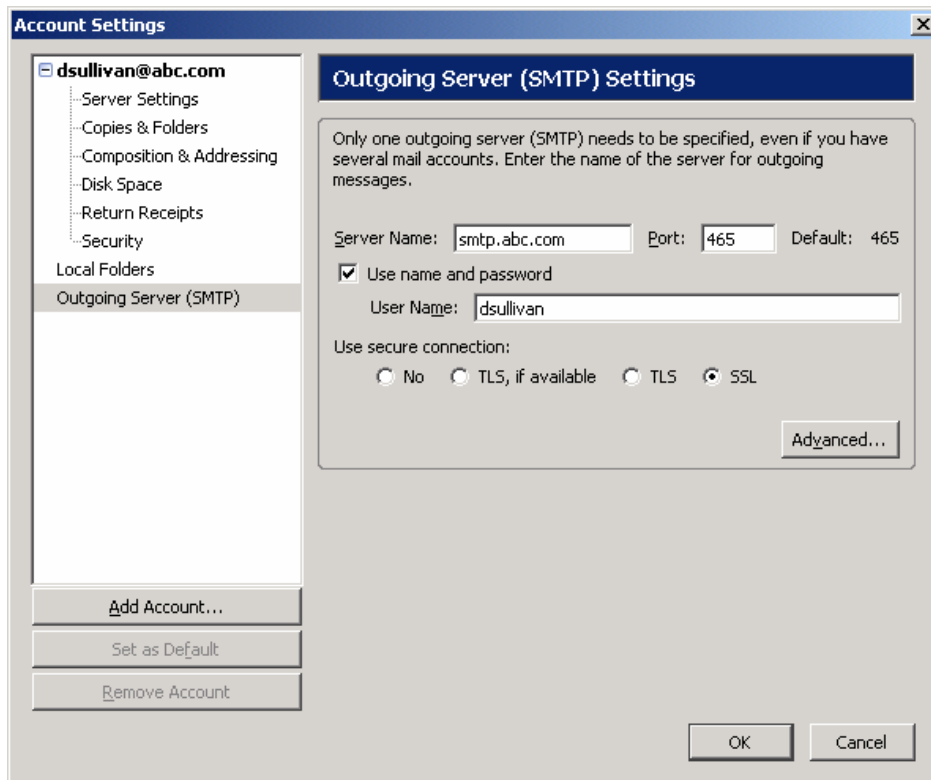


Figure 6.2: Mail servers can be configured to authenticate before allowing client-to-server communications. This example is from the Mozilla Thunderbird email client.


 For more information about SMTP-AUTH, see the specifications at <http://www.faqs.org/rfcs/rfc2554.html>.

Another example of a trust assumption is that servers do not verify the origin of a message. The sending server can put any origin address in the message and send it. The receiving server can then take one of several actions:

- Accept this address as correct and continue to handle the message
- Determine whether the origin address is the same as the machine sending the message; if not, reject the message
- Determine whether the origin address is the same as the machine sending the message; if not, insert the IP address of the sending machine

The first option allows spammers to substitute fake addresses (*spoofing*) and hide the true identity of the sender. The last option, appending the IP address of the sending machine, at least provides some path information if someone were to trace back the origin of a message. Of course, when dynamic IP addressing is used, which is common, one must know both the IP address and the time the message was sent to determine the origin.

Additional protocols have been proposed for authenticating message transmissions between clients and servers as well as between servers and servers. These proposals include DomainKeys Identified Mail (DKIM) and the DKIM Sender Signing Policy.

 For more information about DKIM, see <http://www.ietf.org/internet-drafts/draft-allman-dkim-base-01.txt>. For details about DKIM Sender Signing Policy, see <http://www.ietf.org/internet-drafts/draft-allman-dkim-ssp-01.txt>.

Hiding the Origins of Spam

The cumulative effects of vulnerabilities in the email infrastructure allow spammers to hide behind several methods:

- Faking the sender's address
- Using throwaway email accounts
- Using throwaway domains
- Relaying through third parties
- Using zombies

Authentication methods can help to reduce the problem of fake sender addresses, and increasing the cost of domains can help control the use of throwaway domains. Better security on email servers (for example, eliminating open relays) and other computers can limit the use of third-party relays and zombies.

To summarize, the vulnerabilities in email infrastructure stem from feature-rich but vulnerable clients and inherent weaknesses in SMTP. Vulnerable clients can be a source of email addresses for spammers and the SMTP weaknesses allow spammers to fake origin addresses as well as other header information.

Economics of Spamming

Spamming is a lucrative business. Take Christopher Smith, for example. In May, 2005, United States federal authorities shut down his spamming operation, Xpress Pharmacy Direct, and seized \$1.8 million in luxury cars, two homes, and more than \$1 million in cash from Smith and his associates, according to the Associated Press.

 For more information about the Christopher Smith case, see "Feds: Spamming Made Millions for Dropout" at http://www.sptimes.com/2005/09/12/Technology/Feds_Spamming_made_m.shtml.

Anecdotes such as this highlight the profits to be made in spamming; however, to understand how such a generally unwanted operation can be so successful, one must look into the economics of spamming. The economics of spam can be divided into a number of topics:

- Costs and revenues of spamming
- Distorted costs, or negative externalities, of spamming
- Beneficiaries of spamming

Costs and Revenues of Spamming

Common sense tells us that businesses will continue to operate as long as their revenues are greater than their costs. In addition, the more the revenues exceed the costs, the more likely others will want to get into the business. Economists describe our common sense understandings more precisely with the concepts of supply and demand.

The basic idea is that businesses will provide products and services as long as they can sell their inventory at a price that exceeds the cost of producing the product or service. Simultaneously, customers will continue to buy the products and services as long as the cost does not exceed the perceived value. When prices are low, more customers will come into the market; when prices are high, more producers will come into the market. The logic of the market drives the supply of a product or service to an equilibrium point with the demand (see Figure 6.3). (At least in theory.)





Figure 6.3: *The quantity of a product in the marketplace should balance with the demand for it.*

So if the supply and demand should balance, why is there so much unwanted spam? First, the providers of a product, such as spam, will generate a profit once the fixed cost of their computer equipment, Internet connections, software, and so on plus the cost of sending each additional email message (marginal costs) is paid for. The fixed costs are relatively low—even a desktop PC can generate large volumes of email messages. The cost of generating one more email once the basic equipment and network services are in place is virtually nothing. The result is that the low cost of spam allows spammers to supply large volumes of spam even when the price is low.

A second factor is that even though a spammer might send out 1,000,000 messages, only a small fraction, for example 100 recipients, might respond. Those 100 respondents actually pay the price that covers the cost and profit for the spammer. Although most consumers do not respond to any spam, a small number of consumers responding to even a fraction of all spam they receive can sustain the economic motivations for spamming. According to a 2005 survey by the Pew Internet and American Life Project, 6 percent of respondents claimed to have made purchases in response to unsolicited emails and 13 percent have responded to emails that they later discovered were fraudulent. Such high response rates are not likely to dissuade any would-be spammers from trying their hand at such easy money.

 The full Pew Internet and American Life Project survey report is available at http://www.pewinternet.org/pdfs/PIP_Spam_Ap05.pdf.

The market model works well when buyers and sellers bear all the costs and receive all the benefits of a transaction. Such is not the case with spam.

Negative Externalities of Spam

Economists use the term *externalities* when there are either costs or benefits that are shared by others outside of a transaction. For example, when a steel mill releases emissions into the air, the operation is “free” for the steel mill operators. Those living near the mill pay the price of additional air pollution. The situation is similar with spam. Those who do not want unsolicited email pay the price of having to deal with spam.


Costs Incurred by Others

Businesses and other organizations incur several costs related to spam:

- Wasted bandwidth
- Load on email servers
- Disk and archival storage
- Anti-spam applications
- Employee time

Statistics on the amount of spam are difficult to collect accurately and vary but some estimate that more than 67 percent of all emails are actually spam. This level of network traffic can drive up the cost of bandwidth for both businesses and their ISPs.


Spam also places additional computational and storage load on email servers. With increasing legal precedents and regulations governing electronic communications, some organizations are storing large volumes of emails for extended periods of time. If an organization does not want to risk deleting legitimate email, it might choose to store and archive all emails, including spam. The additional storage and archival costs are born by the recipients of spam, not the senders.

 Legitimate email is also known as “ham” when distinguishing it from unsolicited, unwanted email, spam.

Ideally, spam will never reach a user's inbox. An important element of many IT organizations' strategies for controlling spam includes the use of anti-spam applications. These include both software and network appliances that scan emails before they reach the mail server or, at least, the end user. The cost of acquiring, maintaining, and managing these systems is another cost incurred by those outside of the business transactions of the spam sender.

Another cost that is difficult to quantify is the value of lost employee productivity. Managing spam is annoying, and if too much reaches the inbox, can become a time drain on end users.

As much of the true cost of spam is paid by someone other than the spammer or the person responding to the spam, the economic balance of spamming is therefore distorted.

 You can find a useful anti-spam ROI calculator at <http://www.mcafeesecurity.com/us/products/tools/roi/spam.asp>.

Distorting the Supply/Demand Balance

As depicted in Figure 6.3, free markets will balance themselves so that suppliers will come into the market in just the right number to meet the demand for a product at a particular price. This principal describes the behavior of suppliers in aggregate. Individual suppliers will enter or leave the market depending on the profit they can earn. The suppliers' profits are determined by a combination of their fixed costs, marginal costs, and marginal revenues.

Consider the decision-making processes of a spammer. To get started, the spammer needs a PC, a network connection, and some software. He or she does not have to pay any of the costs incurred by the recipients of spam, so the normal cost information provided by markets is distorted—a fact that has a direct impact on the spammer's decision making.

To maximize profits, suppliers will continue to provide a product as long as the marginal revenue earned is greater than the marginal cost. As Figure 6.4 shows, when the cost to the spammer does not reflect the full cost of spam, the spammer will produce a large quantity of spam (shown as the Distorted Supply line). If the spammer had to incur the full cost of spam, including the costs now incurred by unwilling recipients, the supply would be much less and possibly non-existent (shown as the Undistorted Supply line).

The distortion in the supply of spam is that the market does not convey the proper cost signals to the supplier. The result is the negative impact on email users who are not customers of the spammers. This kind of negative impact, or externality, has occurred before.





Figure 6.4: Without the impact of the true cost of spam on the suppliers, the quantity of spam generated is greater than what the market equilibrium would support.

Correcting for Negative Externalities: Government Regulation

Pollution is a classic example of a negative externality. The market does not have a mechanism to force polluters to incur the cost of pollution. From the perspective of a polluter, there is no cost to the bottom line for polluting and so no reason to stop or limit it. Governments compensate for this type of market inefficiency by imposing a direct cost on polluters through regulations, such as fines for exceeding pollution limits. In an effort to apply this tactic to spammers, many countries have established laws regulating spam, such as the Controlling the Assault of Non-Solicited Pornography and Marketing (CAN-SPAM) Act enacted in the United States in 2003.

 The United States, the European Union, and some of its member states as well as other countries have passed anti-spam legislation. For links to information about these laws, see <http://www.spamlaws.com/>.

The CAN-SPAM Act

The CAN-SPAM Act was passed because spam had moved well beyond being a mere inconvenience or annoyance to a threat to the continued utility of electronic communications. The United States Congress also found that a number of states had passed anti-spam laws but found that the different penalties, standards, and requirements were not effectively addressing the problem.

The CAN-SPAM Act prohibits a number of actions commonly used by spammers:


- Using false or misleading transmission information
- Using deceptive subject headings
- Including false return addresses
- Sending additional commercial mail after recipient objects
- Using address harvesting and dictionary attacks
- Using automated methods to create email accounts for sending spam
- Relaying email through unauthorized access

The law also requires that advertisers provide an opt-out method for recipients as well as a valid physical address of the sender and a clear indication that the message is an advertisement. The legislation includes fines of as much as \$11,000 for violations of each provision. The law also allows for criminal prosecution in cases in which spammers

- Make unauthorized use of another's computer to send spam
- Relay messages through multiple mail servers to hide the messages' origin
- Falsify header information
- Register for multiple email accounts or domain names using false identities
- Falsely represent themselves as owners of IP addresses

Researchers studying the effects of CAN-SPAM have found that the law does help consumers and ISPs block unwanted, unsolicited emails when the senders comply with the law, but technical measures are required when spammers do not comply. Even the effectiveness of technical countermeasures is limited when spammers take actions to evade detection (Source: Matt Bishop, "Spam and the CAN-SPAM Act"

<http://www.ftc.gov/reports/canspam05/bishoprpt.pdf>).

 A summary of the CAN-SPAM Act's requirements for commercial emailers is available from the United States Federal Trade Commission (FTC) at <http://www.ftc.gov/bcp/online/pubs/buspubs/canspam.htm>. The full text of the law is available at http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=108_cong_public_laws&docid=f:publ187.108.pdf. An FTC report on the effectiveness of the act is available at <http://www.ftc.gov/reports/canspam05/051220canspamrpt.pdf>.

Effectiveness of CAN-SPAM

Measuring the effectiveness of CAN-SPAM is difficult. Surveys and research reports do not always agree. According to a report from the United States FTC, since CAN-SPAM went into effect, the spam problem has improved. These improvements are attributable to a number of factors:

- Spam volume has declined, along with consumer frustration, as legitimate mass emailers comply with the law
- Anti-spam technologies have improved and are widely deployed
- The adoption of domain authentication protocols that verify the source of the message is the same as the source listed in the message header

Federal prosecutors are also using the legislation to bring criminal charges against spammers. In the first year, more than 20 cases were brought under the CAN-SPAM Act. In one case, a federal grand jury indicted spammers involved in a large-scale international spam operation that used several techniques to hide the origin of their messages:

- Sending from computers with IP addresses in the Netherlands with domains registered in the Indian Ocean island state of Mauritius
- Using fake addresses in the From line
- Remotely controlling servers in the Netherlands from systems in the United States

In addition, overseas companies were established to disguise their operations and overseas bank accounts were established to facilitate the laundering of profits from the spamming venture (Source: Department of Justice Press Release “Three Defendants Indicted, Fourth Pleads Guilty In Takedown Of Major International Spam Operation,” August 25, 2005—http://www.usdoj.gov/opa/pr/2005/August/05_crm_431.htm).

Although the FTC report argues for definite improvements, the Pew Internet & American Life Project finds mixed success in controlling spam. A 2005 report on spam finds:

- 28 percent of email users with personal email accounts report receiving more spam than a year ago; 22 percent report receiving less
- 21 percent of email users with work email accounts report receiving more spam than a year ago; 16 percent report receiving less

It also reported some improvements, for example:

- 53 percent of email users say spam makes them less trusting of email, an improvement over the 62 percent the prior year
- 22 percent of email users have reduced their use of email due to spam, an improvement over the 29 percent the prior year
- 67 percent of email users claim that spam has made their online experience unpleasant or annoying compared with 77 percent the prior year

 For full details on the Pew Internet & American Life Project report, see http://www.pewinternet.org/PPF/r/155/report_display.asp.

Even with the cooperation of legitimate emailers and the ability of regulators to charge violators, spam continues. Regulators recognize that legislation and prosecution are not enough to stem spamming, and technology will continue to play a central role in efforts to control unsolicited emailing. Before addressing the technical aspects of controlling spam, there is one more element of the economics of spamming that should be addressed.

Cutting Costs and Avoiding Prosecution: The Role of Botnets in Spam

When regulators work with ISPs, together they can trace the source of spam to its origin and find the perpetrators—at least that used to be the case. One can still trace the path of spam through mail servers across the network, but doing so today often leads not to a spammer's server but a compromised system owned by an unwilling participant in a spam operation. Using someone else's computer to distribute spam has two key advantages: it saves the cost of hardware and it makes it more difficult to trace the source of spam.

Spammers have adopted techniques used by malware developers, such as installing malicious programs known as Trojan Horses on victims' computers. The spammer communicates with the Trojan Horse through IRC, FTP, or some other communication protocol, sending instructions and data as needed to direct the distribution of spam. By infecting large numbers of computers, the spammer can create and control a network of "zombies" that perform the bulk of the work for mass mailings. For more technical details on botnets, see John Kristoff's presentation "Botnets" at <http://www.nanog.org/mtg-0410/pdf/kristoff.pdf>.

Beneficiaries of Spam

The obvious beneficiaries of spam include the spammers themselves. When consumers respond to spam, spammers benefit from the profit on any sales transactions but also from the fact that they have the address of a responder. Email addresses of spam responders can command a premium when selling email lists. Unfortunately, this is not the extent of businesses that profit from spam.

In addition to selling products, spammers can make money by generating leads for high-value products such as mortgages. For example, a technology correspondent for MSNBC responded to an unsolicited email with an enticing subject line about low interest mortgages. Within days, he received four inquiries from legitimate mortgage institutions.

The path from the spammer to the lenders passed through a lead generator in the United States who claimed to have purchased the information "from someone else, who in turn bought it from someone else, who in turn bought it from an emailer based in China" (Source: Bob Sullivan, "Who Profits from Spam?" at <http://msnbc.msn.com/id/3078642/>). The lead generator in the United States sold the information to another party who then sold it to yet another party before it was finally sold to a mortgage provider (see Figure 6.5).

When customers complain, legitimate businesses can track down offending lead generators and affiliates—incurring yet another additional cost to businesses directly related to spam. It also requires consumers to take the time to file complaints and provide details so that the business can address the problem. Needless to say, it is often not worth the consumer's time to file the complaint in the first place. The result is that these grey market operations can launder email addresses and lead information providing yet another line of revenue for spammers.

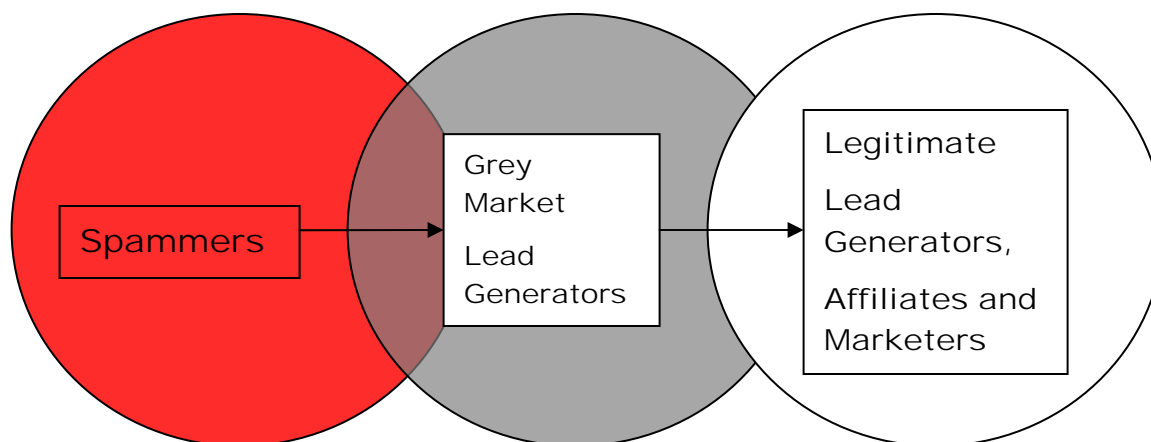


Figure 6.5: Legitimate businesses may ultimately use leads generated from spammers that use a network of grey market lead generators to mask the original source of the leads.

Clearly the economics of spam hold the potential for profit, especially when one is willing to circumvent regulations. Government intervention is certainly helpful—it will at least allow legitimate mass email operators to operate within defined boundaries. Regulations alone will not control spam, and technical measures are required to adequately address the problem.

Spam Management

Just as regulation is not enough to control spam, no single technical measure will completely control spam. The process of identifying and disposing of spam consists of four operations:

- Detection and determination
- Actions in response to spam
- Managing detection methods
- Managing quarantined messages

Together, these tasks constitute the technical aspects of spam management.

Detection and Determination

The first step in spam management processes is to detect suspect messages and determine whether they are actually legitimate messages or unwanted, unsolicited email. Ideally, this process occurs before messages reach the recipient's email server and thus eliminate unnecessary load on the server. The process should operate on all messages; if multiple email servers are in place, the spam detection and determination operations should occur on traffic streams to all of them.

Once email traffic has been intercepted or redirected to a spam detection mechanism, there are several methods for determining whether a message is spam. The common techniques are:

- Integrity analysis
- Heuristic detection
- Content filtering
- Blacklists and whitelists
- Self-tuning
- Bayesian filtering
- DNS block lists

These techniques each have advantages and disadvantages, but together their complementary strengths provide an effective detection mechanism.

Integrity Analysis

Integrity analysis examines the structural characteristics of email messages to determine whether they may be spam. The header, the layout of the messages, and the overall organization of the message can provide clues about the status of a message. For example, a header may have an invalid time zone or a date far earlier than the current date. The body of the message might contain a single line of text in upper case with many whitespace characters and end with an exclamation mark followed by a short paragraph of text and a URL preceded by "Click here." This type of detection is a specific form of the more general process of using heuristics to detect spam.

Heuristic Detection

Heuristics, or rules of thumb, are often used to craft detection rules. Rules are typically of the form:

```
If <some property of the email>  
Then spam score = spam_score + <confidence factor>
```

The condition, "some property of the email," can be any pattern that is indicative of spam. For example, the presence of a spam-tool name in the header or the use of upper-case letters in the subject line or body of the message.

The confidence factor is a measure of how well the pattern in the condition indicates spam. Some patterns are very often associated with spam, such as a URL with the word “remove” in it and a subject line that contains a unique identifier. These would qualify as high-confidence factors. Other patterns often found in spam, such as color-coded HTML text is also found in legitimate email, so the confidence factor would be lower.

The spam score is the sum of confidence factors of all rules that are true for the message. When a message’s spam score exceeds a predefined threshold, the message is categorized as spam.

One advantage of heuristic rules is that they allow for custom coding of filters that can take into account the location of a pattern (for example, in the header versus the body) and thus can make finer distinctions than just looking for a particular pattern of characters anywhere in a message. Another advantage is that different combinations of rules can detect a wide variety of spam. One spam message might be detected because of the use of colors and keywords while another is detected because of an invalid date header and a suspicious phrase such as “As seen on national TV!” At the same time, the confidence factors and thresholds can be set to ensure that no single rule can trigger the classification of a message as spam. They can also be adjusted to minimize the likelihood of identifying legitimate mail as spam.

Perhaps the biggest disadvantage of heuristic rules is the time it takes to develop them. Designers have to take into account how the condition in the rule might detect legitimate email and how the rule interacts with other rules. Unlike virus detection, which can use a single signature to identify a particular piece of malware, spam detection is done with a set of rules so that full rule sets must be developed and tested before they are released.


Content Filtering

Content filtering uses lists of words and phrases that indicate spam or offensive material that is banned within an organization. Content filtering is an efficient technique for identifying blatantly offensive messages but it lacks the ability to distinguish legitimate uses of a term from inappropriate uses. Content filtering lists are generally limited so as not to mistakenly categorize legitimate email based on the use of terms that have both appropriate and inappropriate uses in a business context.

A specialized type of content filtering involves domain name reputation technology. The domain name reputation technique examines the URLs within each email message and blocks those messages that contain links to malicious or spammer Web sites. The technique works to block spam, phishing messages, and messages containing links to malicious payloads such as spyware. The accuracy of this technique can be very high but is dependent on having reliable and up-to-date lists of suspicious domains. Such lists are very costly to maintain and are typically beyond the reach of all but the largest and best-funded anti-spam vendors.

Blacklists and Whitelists

Blacklists and whitelists are essentially lists of known spammers and known non-spammers, respectively. The advantage of these lists is that they allow the spam management software to quickly categorize messages based on information in the header. More computationally intensive operations, such as applying heuristic rules and calculating spam scores, are not necessary.

 The Open Directory Project (<http://dmoz.org>) maintains a list of blacklist providers at <http://dmoz.org/Computers/Internet/Abuse/Spam/Blacklists/>.

Self-Tuning

The techniques discussed so far are generalized to apply to all email users. For example, if an email user is categorized as a spammer on a blacklist, that user is considered a spammer for everyone. With self-tuning techniques, anti-spam applications can adjust their categorization rules to take into account the type of email individual users or organizations receive. When a message is sent from a user with a history of sending legitimate emails, the spam score assigned by heuristic rules can be adjusted based on patterns derived by self-tuning.

Bayesian Filtering

Bayesian filtering is a mathematical method for using the probabilities of a word, phrase, or symbol being found in a spam message. Unlike content filtering, which is based on a list of commonly used words and phrases compiled by a human, Bayesian filtering methods analyze large numbers of spam and legitimate mail to calculate precise measures of a word's likelihood of being found in spam.

Without going into too many details of the math behind Bayesian filtering, we can still outline the basic points. First, Bayesian filtering uses probabilities that are measured between 0 and 1 that indicates the likelihood of a particular hypothesis; 0 indicates with certainty that the hypothesis is false, 1 indicates with certainty that the hypothesis is true. If a weather forecaster says there is a 0.5 probability of rain, it is just as likely to rain as not to rain. Probabilities calculated for Bayesian filtering are different from scores used in heuristic rules. Scores are intuitive measures assigned to rules by individuals; they may be adjusted to improve accuracy when used with other rules. Probabilities are calculated using a specific formula.

Second, Bayesian filtering uses information from both example spam and legitimate mail. Consider the word "free." It is used in both spam and legitimate email. In a sample of 10,000 legitimate emails, "free" may appear 200 times; the same word might appear 1000 times in a sample of 10,000 spam messages. Clearly the word "free" is a good indicator of spam, but how good? How should you weight your belief that an email with that word is actually spam? A Bayesian formula, known as Naïve-Bayes, can tell you the answer (see Figure 6.6).





Figure 6.6: Bayesian filtering takes into account the frequency with which words appear in both spam and legitimate emails.


Another factor of Bayesian filtering is that it can adjust and adapt as the number of examples of spam and legitimate email grows. For example, email sent by a user as well as email sent from email users on an organization's whitelist can be considered illustrations of legitimate email. Messages identified as spam by email users or deleted without opening may be considered spam and used to adjust the spam probabilities of words occurring in those messages.

 For more details about the math behind Bayesian filtering, see http://en.wikipedia.org/wiki/Naive_Bayes_classifier.

Spammers try to avoid detection using several techniques:

- Using long stories that skew the statistical measures and lessen the overall probability measure that a message is spam
- Inserting randomly selected words from a dictionary to again skew the spam measure
- Adding unrelated information, such as the text from a news story, into the message.

Fortunately, these techniques that can fool a Bayesian filter can be detected by other spam-detection techniques such as heuristic rules. Spammers may be able to get around one anti-spam measure, but it is difficult to fool multiple techniques simultaneously.

 For a discussion on the limits of Bayesian filtering, see William S. Yerazunis' "The Spam-Filtering Accuracy and How to Get Past It" at http://crm114.sourceforge.net/Plateau_Paper.pdf.

DNS Block Lists

DNS block lists use the IP address of the sender to identify known spammers. DNS blocks lists are publicly available on the Internet so that email administrators or anti-spam vendors do not need to maintain individual lists. Of course, like other commonly maintained resources, the quality can vary from poor to good. It is not uncommon for DNS block lists to contain the addresses of non-spammers (false positives).

A combination of detection techniques, from examining the structure of a message to analyzing word frequencies, can provide highly accurate means of detecting spam without generating unacceptable numbers of false positives. Of course, once spam has been identified, the anti-spam system must decide what to do with it.

Actions in Response to Spam

Once a message has been categorized as spam, there are generally three actions the anti-spam system can take:

- Delete the message automatically
- Quarantine the message
- Tag the message as spam and deliver it

The appropriate action will depend on user preferences and the level of confidence that the message is truly spam.

When a message is deemed to be spam with certainty—for example, the message is from an address on a blacklist, the message may be deleted without user intervention. Addresses included on blacklists are used only when all email from those sources should be blocked. Email administrators may also configure an anti-spam system to automatically delete any message with a spam score above a particular threshold. The drawback of deleting messages is that in the case of a false positive, the recipient has no way to know the message ever existed or to review it before disposing of it.

Quarantining is an alternative to deleting messages in which messages are isolated. The recipient's inbox is not cluttered with spam, but messages are not permanently deleted either. Typically, the messages are preserved in an isolated folder and a list of quarantined messages is sent to the recipient. The recipient can retrieve any of the quarantined messages or delete them if they are actually spam. Messages in quarantine are usually purged after a predefined period of time (such as 30 days) if no other action is taken on them by the recipient.

The third option is to tag an email message and send it on to the recipient. For example, the prefix spam could be added to the subject line to highlight the fact that the message is likely spam.


The appropriate action will depend on a combination of factors, including tolerance for false positives and the amount of spam received. Low tolerance for false positives calls for quarantining or tagging. Receiving large volumes of spam argues for automatic deleting. When both factors are in effect, users and administrators must find their own suitable balance.

Blacklist and Whitelist Management

As previously noted, blacklists are used to define addresses of known spammers and whitelists are used to define addresses of known legitimate emailers. These lists are not static and can change frequently as the needs of the organization and individual email users, as well as the behaviors of spammers, change. These lists must be updated, for example, when

- Spammers change domains and begin sending spam from a new source
- Businesses acquire new customers or business partners
- Email users find themselves on mailing lists they do not want to be on
- Legitimate emailers have been added to blacklists by mistake

Both blacklists and whitelists should be managed globally and individually. Global lists apply to all messages sent to a server. Email administrators are responsible for maintaining these. As it is difficult, if not impossible, for a single organization to track all known spammers, many email administrators use publicly available real-time blacklists.

 The Spamhaus Block List (<http://www.spamhaus.org/sbl/index.lasso>) and SORBOS (<http://www.nl.sorbs.net/>) are two popular blacklists. For a comparison report on a number of blacklist sources, see the latest report at http://www.sdsc.edu/~jeff/spam/Blacklists_Compared.html.

Anti-spam systems should also allow individual users to configure personalized blacklists and whitelists. Server-level blacklists should catch many spammers, so the personalized blacklist can target lower-volume spammers and phishers. In addition to managing blacklists and whitelists, email administrators and users also need to manage the quarantine process.

Quarantine Access, Search, and Management

Quarantining messages can help to counter the effects of false positives but it does impose additional management tasks on administrators and users. Administrators have to take into consideration a number of factors, including:

- How much disk space to allocate to quarantine folders
- How long to retain quarantined messages before deleting them
- How to adjust retention policies for different types of users

For users, the issues tend to center around access and search. How often is the user notified about quarantined messages and how is the information delivered? A single message once a day listing all quarantined messages works well in many cases. Users may also want immediate access to the quarantine folder through their email client. With long retention periods, the number of spam messages in a quarantine folder will grow and users will need search tools to sort and filter quarantined messages as they look for legitimate messages that may have been incorrectly categorized as spam. The ability to manage quarantined messages is just one of the considerations one should take into account when evaluating anti-spam systems.

Evaluating Anti-Spam Systems

There are a number of anti-spam systems available today, both from vendors and open source projects. Choosing among these options can be a difficult task, so it is important to focus on key features:

- Catch rate
- False positive rate
- Manageability and reporting

In addition to these, there are, of course, the ever-present concerns about integration and reliability that come with any enterprise-scale application.

Catch Rates

An anti-spam system's catch rate is a measure of how well it detects spam. As noted earlier, spammers can often trick one detection method at a time but it is difficult to fool multiple methods simultaneously. For high catch rates, consider systems that use multiple techniques, especially Bayesian filtering, heuristic rules, and domain name reputation technology.

False Positives

False positives, or legitimate email categorized as spam, are a serious problem for anti-spam systems. It is generally preferable to allow some spam through rather than risk blocking legitimate email. False positive rates can be minimized with the use of whitelists, heuristic rules, and self-tuning. Systems that allow individual users to customize whitelists and train Bayesian filters can also help reduce the chances of generating false positives.

Manageability and Reporting

Manageability is one of those general characteristics exhibited by well-designed systems. In the case of anti-spam systems, there are three areas administrators should consider:

- Flexibility in configuration for different sets of users
- Ease of updating blacklists and whitelists
- Ease of quarantine access and management

Different sets of users will require different policies governing their use of email. For example, some departments, such as legal affairs or human resources, may have no tolerance for false positives. They may need all suspect messages with high spam scores quarantined and those with moderately high scores labeled and delivered to the recipient. Other accounts, such as a general customer service email account, may receive a great deal of spam because the email address is published on the company Web site. In that case, messages with moderate or high spam scores may be automatically deleted while low scoring messages are quarantined for short periods of time, such as a day or two, before being deleted.

Email administrators should be able to edit their whitelists easily. Email users should also have the ability to specify custom lists based on their own email patterns.

Finally, both administrators and users should be able to review quarantined messages, transfer them to inboxes, and delete them as necessary. Additional searching, sorting, and filtering features will help when large numbers of messages are quarantined.

Anti-spam systems should also provide administrative reports that allow managers to track the volume of email messages analyzed, the number of spam messages detected, the volume of storage used for quarantine, and other key indicators of the performance of the system.

Spam management depends upon technical and regulatory measures. Although regulations help to define the boundaries of appropriate mass emailing and allow legitimate marketers to stay within the law, they cannot prevent determined spammers from flooding inboxes with unsolicited, unwanted email. Technical solutions, especially multi-tiered methods for spam detection, are the key to controlling the impact of spam on email infrastructure and end users.

Summary

Spam is nothing new. One of the earliest examples of spam dates back to the late 1970s, and by the 1990s, spamming email and listserv systems was growing in automation and sophistication. Today, spammers exploit vulnerabilities in email infrastructures, leverage compromised hosts (“zombies”), and adapt their messages to avoid detection from a number of anti-spam techniques. Although spamming is illegal in many countries, the problem continues. The economic benefits of spam are based on the fact that even very low response rates can generate enough revenue to more than cover the direct costs to spammers. This, of course, does not cover all the costs, because, like pollution, the cost of spam is shared by many, not just those that create the problem.

Anti-spam systems employ effective and highly accurate methods for detecting spam. A combination of several methods provides the best defense against the adaptive nature of spammers. In addition to accuracy, manageability and reporting are key considerations when selecting an appropriate anti-spam system for an organization.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.