Global Information Assurance Certification Paper

Did You Get My Email?

A Look At Spam Filtering And How To Avoid Getting Your Innocent Email Caught In Its Trap

Ray Ellington
July 20, 2004
GSEC Practical Assignment V.1.4b

Abstract
With the recent increase in the amount of spam being circulated on the internet, more and more organizations are placing spam filters at their gateways and on their client machines. With the increase in spam filters comes an increase in the likelihood that innocent email will be falsely identified as spam. End users of spam filters need to understand their products and tune them as necessary. We, as users of email, must be aware of how our mail may be treated by spam filters while we are composing it. Systems Administrators must be aware of the restrictions that large ISPs like AOL and Yahoo are imposing on sending domains. Being a network administrator who has felt the pain of being blocked by large ISPs has prompted me to research the subject of technologies that are working toward eliminating false positives in the mine field of spam filters.

I am going to start off by discussing some common methods used by modern spam filters for classifying email as spam. I will then go on to discuss some new technologies being developed to guarantee that email will reach the person it is supposed to. In closing, I will explain some best practices that administrators can use to better ensure that email from their users gets to the person it needs to.

Introduction
Has your email ever been classified as spam when it really wasn't spam? Would you even know if it was? Did the short email you sent your boss explaining that you wouldn't be able to come into work on Monday make it to him or was it deleted as spam? The problem gets worse as important bulk mailings from organizations and universities to their students, clients, and customers are falsely classified as spam and never even reach them. False positives cost US businesses around 3.5 billion this year. An estimate for the amount of time US employees spend trying to determine if their email reached the intended recipient is $50.00 per year per employee.[1] According to Ray Everett-Church, counsel for the Coalition Against Unsolicited Commercial Email, between 5-15% of opt-in email gets blocked due to the inaccurate blacklists used by large Internet Service Providers.[2]

Preventing your email from being classified as spam begins with understanding the functions that modern spam filters use to determine whether an email is spam or not. This methodology, in my opinion, is similar to learning how to defend against hackers by understanding their methods of attack.

Common Methods of Spam Filtering
Spam Filters have been getting better as more effort goes into developing new techniques. The rising cost of spam has contributed to an intense focus on new spam filtering technologies over the past few years and more companies are packaging their products with some sort of spam filtering capabilities. Microsoft Outlook 2003 has some very effective spam filtering built into it as does Microsoft Exchange 2003. AOL, Earthlink, Hotmail, and other web enabled mail clients

now have spam controls for the end user. The methods being offered to the end user have improved from simple keyword based filtering to filters that actually learn the difference between valid email and spam through a training process.

One technique of spam filtering at the gateway that has been around for quite some time is known as Realtime Black Lists. Realtime Black Lists (RBLs) are lists of IP addresses or domain names that spam has been known to originate from. The IP addresses are collected and maintained by many different organizations. Spam filters can perform lookups against the RBLs to determine whether the IP address initiating the connection is on the list or not. This lookup is performed by reversing the octets of the IP address and querying under the RBL's domain. For example (using spamhaus.org ads), if a sender's IP address is 10.1.1.5, the spam filter would attempt to resolve the address resource record for 5.1.1.10.sbl.spamhaus.org. The name server for sbl.spamhaus.org would normally return 127.0.0.2 if the server is listed on their Realtime Black List. If the server is not listed on the RBL then the name server will return that the server name is invalid.

Organizations can create their own black lists which is a very time consuming process or they may also rely on other organizations who are dedicated to keeping track of well known spamming addresses. Some of the largest organizations that do this are the following:

- Spamhaus.org www.spamhaus.org
- Ordb.org www.ordb.org – List of mail servers that are known to be open relays. Open relays are used by spam senders to anonymously send email.
- Dsbl.org www.dsbl.org – Distributed Server Boycott List
- Spamcop.net www.spamcop.net

Another method of spam filtering is called Content Filtering. Content Filtering is a crude form of filtering which uses lists of keywords known to be contained in spam. The lists are normally created by the email administrator or the end user of the email client software. When an email is received, the header, subject or body of the email is compared against the list of key words. If the keyword is identified then the mail is considered spam and either deleted or moved into quarantine.

A smarter method of keyword type filtering is called Heuristic Filtering. Heuristic Filtering performs an analysis on the email and then compares the email against a database generated by "training" the filter. Training the filter consists of providing a number of spam as well as non-spam messages to the filter so that it can learn patterns of keywords which may be considered legitimate or non-legitimate for your organization. A score is then assigned to the email based on the heuristic analysis. Email which falls into the values (usually set by the user) considered to be spam is then either placed into quarantine or deleted. Spambayes and Outlook 2003 use this technique of spam filtering.

Some spam filters allow the user to adjust the rate at which incoming email is received from one IP address. Spammers may target a mail server with a bulk mailing destined for each user in the organization. Rate control mechanisms will pick up on this and prevent further mailings from the sender. Rate control can be specified per the following:

- Number of sessions per time unit
- SMTP session timeout
- Number of messages sent per SMTP session

A brutal method of obtaining valid email addresses and sending spam is referred to as a Dictionary Attack. A Dictionary Attack is performed by the spammer attempting to send mail to a list of common usernames proceeded by the domain address of the organization hoping that one of the combinations will match an email address used in the organization. For example: bob@domain.com, bob1@domain.com, admin@domain.com, webmaster@domain.com, roger@domain.com, etc. If only one out of hundreds of attempts is successful, the spammer will have done their job. Even worse, the email may contain a piece of code that allows the spammer to know if it reached a valid recipient. To protect against this type of attack, some spam filters will query an LDAP server or another database which contains a list of the true email addresses used in the organization before accepting email from that sender. For example, the Barracuda Spam Firewall will, through a read-only connection to the Active Directory Database, know which addresses are valid for the organization and only accept mail for those addresses.

Domain spoofing is a popular method of sending spam. Domain spoofing involves using someone else's domain name when sending an email. Anti domain spoofing methods are growing in popularity at becoming a leading method to cut down on spam. It is used widely by spammers and worms to lure people into divulging sensitive information. For example, phishing schemes which send email to users as coming from support@paypal.com asking the user to verify their account. The person receiving the email may believe it to be true since it is coming from support@paypal.com when actually it is a spoofed domain. Mass mailing worms would be prevented from sending mail from compromised machines if this technology were in place because most of the worms will use spoofed From: addresses when sending mail. If a receiving mail server sees mail coming from the domain paypal.com and the IP address 192.168.1.2, it will know that paypal.com does not send mail from 192.168.1.2 and therefore close the connection.

Organizations pushing these technologies believe that this will cut down on a large percentage of spam since most of it is sent from open relays and compromised machines. Most of the technologies are still in development at this time. The three leading technologies are as follows:

1. SPF Protocol – SPF stands for Sender Policy Framework. Mail Transfer Agents can use SPF to verify the envelope sender (SMTP MAIL FROM or return-path) address during SMTP time. If the address does not match what is published by the domain owner, the mail is rejected. SPF works through the domain owner publishing a list of mail servers that are authorized to send mail in its domain records using the TXT record type. When a mail server receives an email, it checks the domain in the return-path or the "From:" header then performs a lookup against the SPF record of the domain in question. If the IP address is not in the record, the mail is rejected as spam[3].

2. Microsoft's Caller ID – Still in development and recently merged with the SPF protocol and renamed Sender ID. The concept is similar in nature to SPF but is different in the way the way the records are published and the content of the mail that is analyzed for spoofing. Caller ID filters can analyze more content of the email than the SPF protocol does[4]. Domain owners would publish a list of their outgoing Mail Transfer Agents in their DNS zone using XML. The receiving mail server would analyze the message to determine the sender IP address the compare against the list of approved Mail Transfer Agents published by the domain owner.

3. DomainKeys – This technology was developed by Yahoo! Inc. It uses Public/Private Key Cryptography to generate a signature which is injected into the email header. The receiving mail server checks the signature against the sender's public key which is made available via DNS[5].

Guaranteed Delivery Methods

The increasing cost of false positives as well as the frustration of sending email only to have it fall into the spam category is pushing the need for a way to guarantee that your email will arrive at the inbox of the intended recipient. One of the companies listed below, Bonded Sender, quotes the following from the Washington Post on their website:

"AOL's spam Filters blocked nearly 100 notices that Harvard University had emailed to applicants telling them whether they had been accepted or not"

At my place of employment we have performed mass mailings to candidates informing them that they could now come online to choose a test center within their city limits for a very important exam pertaining to their career. Some of the candidates did not receive the email due to being blocked by spam filters. Those candidates, due to late notice, were faced with the possibility of having to travel to another city for testing versus testing within an hour of their residence.

Our organization's sending IP address was recently put on a 24 hour temporary block by America Online. AOL has made it very easy to allow users to report

spammers by giving them access to a “Report Spam” button. Apparently, even after they have gone through a double opt-in process to get on our mailing list, some users forgot about that and reported our mail as spam. This cut off all help desk communication with anyone having an AOL address.

Non marketers are beginning to lose their trust in email as being a reliable source of communication and marketers are reconsidering their use of email as a marketing strategy. Below are some companies that are trying to fix the problem of false positives and restore trust in email.

A company called Habeas has created a method called Sender Warranted Email. To use this method, a company must first qualify by meeting a set of standards which are listed below. Once the company has shown that they meet the standards, they are provided with a special set of copyrighted x-headers that they insert into their outgoing mail. Participating anti-spam software providers and ISPs will perform a check on incoming mail. If the mail contains the Habeas header, it will be allowed to pass through the organization’s spam filters.

For an organization to use the Habeas technique of guaranteed email delivery, email must first pass standards to be considered Anti-Spam. These include the following:

1. Offer an Unsubscribe function on all email
2. Have a removal policy for bounced emails such as an average bounce rate of no more than 5% bounced mails for each mailing.
3. Obtain verified permission from the recipient to receive emails from the sender. The permission can consist of the following:
    a. One-to-One: The consumer has a pre-existing business relationship with the sending party.
    b. Confirmed Opt-in: Consumer explicitly agrees to have their email address placed on a mailing list.

Once standards have been passed, the sender is allowed to use a special header which is seen by participating ISPs and Anti-Spam software developers. Then the mail is allowed through the spam filters.

The Habeas x-headers are interesting in that they are very small three-line haiku poems. Titles and phrases can be protected by trademark law, but are too short to be considered literary works under copyright law. This allows Habeas to both copyright and trademark their headers. Any spammer who forges the Habeas x-headers to make it through an organization’s spam filters can be sued for both copyright and trademark violation. To this date, Habeas has already shut down spammers in successful court actions[6].

Another big push in the guaranteed delivery of email is coming from Microsoft and other large corporations. Microsoft especially has recently been pushing a

program named the Bonded Sender Program by a company called Ironport. The Bonded Sender Program requires an organization to first meet a strict set of requirements to become a bonded sender. Once the requirements have been met, the sender pays an annual license fee and posts a bond which varies in amount due to the volume of email the organization sends per month. They are then added to the Bonded Sender Whitelist which is used by large ISPs to determine if the email can pass through their spam filters or not. The ISPs install a software component on their mail servers which performs a lookup against the Bonded Sender Whitelist for the sending IP address when the SMTP connection is established. The lookup is performed using the same method as mentioned earlier when discussing Realtime Black List lookups. The bonded organization is allowed a certain number of complaints from recipients per month. If that rate is exceeded, portions of the bond are deducted from the organizations bond.

To have an IP address bonded by the Bonded Sender Program, the sender must be approved through an application process. The process consists of meeting a set of standards in the following areas[7]:

1. Accountability – Participating senders must ensure that the mail infrastructure used to send Email Messages is well maintained and operated in a responsible manner.
2. Transparency - Participating Senders must ensure that Email Messages are truthful and accurately identify the source of the message.
3. Security - Participating Senders must ensure that there is reasonable security for networks used to send Email Messages and store recipient information.
4. Disclosure - Participating Senders must ensure that the following is clearly and conspicuously disclosed at the point of collection of email address and Related Personal Information (a link to a privacy statement is insufficient).
5. Consent - Participating Senders must ensure that consent with appropriate disclosure or a prior business relationship exists prior to sending Commercial or Promotional Email Messages.
6. Unsubscribe - Participating Senders must ensure that the Recipient's requests to discontinue receipt of Commercial or Promotional Email Messages are honored.
7. Responsiveness - Participating Senders must ensure that all parties involved in the sending of Email Messages cooperate with the program administrator to resolve any issues regarding Program Requirements by responding in 5 business days of notice, and by taking corrective action within 15 business days of notice.

A sender may also be certified at a slightly higher level by Bonded Sender, which they refer to as the Bonded Sender Plus standard. To be certified under the Plus standard, senders cannot rely on a prior business relationship as consent for sending email and senders are not permitted to rent, share, purchase, or sell

email addresses to third parties. Also, double opt-in is the only method in which a sender can be permitted to send bonded email to the recipient. Double opt-in consists of the recipient requesting to be added to an organization's mailing list, the organization sending a confirmation email to the recipient, and the recipient responding to the confirmation email by visiting a link or replying.

Best Practices

There are some actions a sender or mail administrator can take to ensure that their mail will reach the intended recipient without being classified as spam. Many of these are required if you want to enroll into a guaranteed sender program such as Bonded Sender or Habeas.

1. Do not send mail from a dynamic IP address.

Be sure that you are not sending email from a dynamic IP address. Some of the larger ISPs will refuse mail sent from a dynamic IP address to avoid spam sent from home users whose machines have been compromised by spam sending worms.

2. Be sure to have a reverse lookup record for your sending mail server.

RFC1912 Section 2.1[9] states that you should have a reverse record for all of your mail servers. When you send email your mail server tells the receiving server its IP address. The receiving server may perform a reverse lookup on that IP address to determine which domain it is assigned to. The receiving server may then compare the domain name which the IP address is assigned to against the domain name received from the "HELO" command to verify that they match. If they do not match, mail may be refused from your server.

3. Maintain List Cleanliness

Keep mailing lists up to date. If an ISP sees a large number of "unknown senders" originating from a single domain then they may begin to classify your mail as spam. This is referred to as your "bounce rate". Try to keep a low bounce rate by removing bad addresses from your mailing lists.

4. Avoid a large number of recipients in your headers

If you are going to perform mass mailings, be sure that you send each mail individually. Email sent to large ISPs with a large number of recipients in the "To:", "CC:", or "BCC:" fields may be stopped by spam filters.

5. Avoid suspicious words in the subject or content of the email

Be careful of using misspelled words, capital letters, or a large number of special characters in the subject or content of your email. Avoid words which may trigger spam filters such as "sale", "free", "mortgage" or "offer". Also, try not to use colored or non-standard fonts too frequently in your email. Avoid sending email consisting of a single inline image.

6. Whitelists
Request that people you wish to send email to add you to their whitelists. This can be done at the user level using your domain name or at the server level using either domain name or IP address. Once added to a user's whitelist, anything that may have triggered your mail to be spam will be ignored and delivered.

7. Use a spam filter for testing
If you are concerned about an email being treated as spam, you may want to first send the email through a spam filter to see how it is treated. Some spam filters will assign a score to the email and show you why it was scored as it was. You can alter your email according to the traits identified by your spam filter then resend until the score becomes lower.

8. Avoid having open relays
Mail servers which act as open relays are destined to be placed on the blacklist maintained at ordb.org. Open relays allow spammers to send email to their targets anonymously through the user of your mail server. To test if your mail server is an open relay, go to ordb.org and perform the "test an open relay" check.

I see email going the way of the default firewall policies on older firewalls. I remember when many firewalls were shipped with a wide open default policy which left it up to the user to explicitly set *deny* rules to block traffic they did not want into or out of their network. Now, firewalls are mainly shipped with closed policies where the user begins building the Access Control List specifying what type of traffic they would like to explicitly *allow* into their network.

This is similar to what is happening with current email policies. For example, organizations are placing spam filters on their gateways and client desktops which are blocking more mail than they should. Additionally, users are being given the option of blocking mail from anyone who is not in their address book. Even America Online is giving users the option of blocking mail from anyone but other AOL members.

If you want to get mail through to the recipient, your best option is asking them to put your address on their white list. Firewalls manufacturers and security administrators did the right thing when they began setting their ACLs to implicitly deny rather than implicitly allow. Maybe it will be a good thing for email to follow this development in the long run. Only time will tell.

References

1) Sturgeon, Will. “Anti-spam efforts more costly than spam?” 13 August 2003.URL:http://www.silicon.com/research/specialreports/thespamreport/0,39025001,10005575,00.htm

2) Claburn, Thomas. “Microsoft Throws Its Weight Behind E-Mail-Accreditation Program” 5 May 2004. URL:http://www.securitypipeline.com/showArticle.jhtml?articleID=19502271

3) Weng Wong, Meng and Lentczner, Mark. “Sender Policy Framework (SPF) – A Convention to Describe Hosts Authorized to Send SMTP Traffic” May 2004. URL: http://www.ietf.org/internet-drafts/draft-mengwong-spf-01.txt

4) Robichaux, Paul. “The Sender Policy Framework and Caller ID for Email” 11 March 2004. URL:http://www.winnetmag.com/Windows/Article/ArticleID/42136/42136.html

5) DomainKeys: Proving and Protecting Email Sender Identity http://antispam.yahoo.com/domainkeys

6) Cox, Tim. “HABEAS WINS COURT VICTORY OVER SPAMMER WILLIAM CARSON” 6 April 2004 URL: http://www.habeas.com/pr15.html

7) Bonded Sender Program - Email Standards URL:http://www.bondedsender.com/standards.jsp

8) Poteet, Jeremy. Canning Spam: You’ve Got Mail (That You Don’t Want) Pearson Education, 13 May 2004

9) RFC 1912 – Common DNS Operational and Configuration Errors URL:http://www.faqs.org/rfcs/rfc1912.html

10) Barracuda Networks URL: http://www.barracudanetworks.com/