NISCC Technical Note 02/04
Issued 27 February 2004

Spam Mitigation Techniques

Key Points

- **Spam is an ever increasing problem, which is costly in itself and may also be used as a delivery method for a range of electronic attacks**
- **Address list collection and selling is increasingly sophisticated**
- **Insecure Internet servers are widely used to distribute spam**
- **There are measures users can take to reduce the risk of becoming a target for spam**
- **Various types of Spam filters can be used, with different characteristics**

National Infrastructure
Security Co-Ordination Centre
PO Box 832
London
SW1P 1BG

Tel: 020 7821 1330 Ext 4511
Fax: 020 7821 1686
Email: enquiries@niscc.gov.uk
Web: www.niscc.gov.uk

Introduction

1. The objective of this technical note is to make the system administrators in the NISCC constituency aware of the various techniques available to minimize the ever increasing amount of unsolicited email (generally advertising) referred to as "spam". Spam email is often crafted to protect the identity and location of the sender, and will often be delivered through a variety of deception techniques designed to get around current filtering technologies. Spam may take many forms such as:

- Chain letters
- Unsolicited petitions, surveys and advertisements
- Mass-mailers and other malware that forge header information
- Bank scams
- Lottery frauds
- Offensive material

2. This paper also provides a description of the various techniques of spam mitigation and gives a list of best practices to aid computer administrators and users. The techniques and tools outlined in this document have not been evaluated by the NISCC staff. They are listed to highlight the broad spectrum of spam filtering methods and to describe how they are designed to function.

Background

3. Spam is one of the unavoidable pitfalls of being connected to the Internet. It is very difficult to prevent because each person connected to the Internet has the option to spam all other users that have a valid email address. Consequently, the number and variety of spammers is sufficient to have caused an increasing public outcry, that has in some instances lead to anti-spam legislation. For example, it is a criminal offence in the United Kingdom to send spam to private citizens. Companies have to get permission from an individual before they can send them an email or text message. Under the new law, spammers could be fined £5,000 in a magistrates court or an unlimited penalty from a jury, but they would not be sent to jail. For more information on anti-spam legislation see:

- http://news.bbc.co.uk/1/hi/technology/3120628.stm
- http://www.spamlaws.com/state/summary.html
- http://www.computerworld.com/governmenttopics/government/legislation/story/0,10801,86383,00.html
- http://www.spamlaws.com/eu.html
- http://www.spamlaws.com/docs/2002-58-ec.pdf
- http://icocms.amaze.co.uk/DocumentUploads/New%20rules%20on%20email%20marketing.pdf
- http://www.informationcommissioner.gov.uk/eventual.aspx?id=783

4. Spam poses the following threats to the business environment:

- Networks of limited bandwidth become congested with the volume of spam and the timely delivery of legitimate email suffers
- The spam email may be accompanied by viruses or other malware
- Damage to employee morale as a result of fear of being punished as merely a passive recipient of offensive or potentially illegal material

- An employee could initiate a law suit against the employer if the employer does not take adequate steps to protect the employee from the delivery of offensive material
- Lost productivity on the part of the recipient who needs to manage the influx of bogus email

5. As an example of lost productivity consider an organisation where 2000 employees use email daily, where each employee spends 5 minutes each day eliminating spam from their inbox. Each day this organisation loses just under 167 man hours, or 833 man hours each week, or 43,300 man hours per year in time that could otherwise be spent on the business of the office. Suppose the average hourly rate of these employees is £10. In this scenario this organisation is spending £433,000 annually for their employees to manage the influx of spam. Open source estimates vary widely on the issue of lost productivity; however it is mutually agreed that the problem of managing spam is increasing and that measures to control the influx of this type of junk mail must be adopted.

6. Understanding the nature of a spam attack will yield a better understanding of the methods of defence. Therefore, this paper offers a brief overview of several of the numerous methods whereby spammers find their way to a victim's computer. They are:

- Dictionary attacks – use randomly generated addresses
- Poor user practices – unwittingly identify an address as valid
- Purchased address lists – spammers sell their address lists
- Automated address harvesters – address collection software
- Open mail server – mass delivery mechanism
- Open web proxy server – mass delivery mechanism
- Other proxies – mass delivery mechanism
- Malicious code – used to create open relays
- Hijacking registered but unused network addresses - border gateway protocol route announcement

Dictionary Attack

7. Generally, dictionary attack programs are designed to generate a brute force assault that will try to guess a password or key by exhausting all possible combinations from a precompiled list of values. Spammers commonly use this automated attack technique to randomly generate email addresses instead of trying to guess passwords. In a sense, randomly guessing a valid email address is the "password" for the spammer, as they use it to target their victims. This approach is successful because email addresses typically contain some form of a user name such as "Robert@domain.com" etc. Consequently, they are easily guessed by the attack program. It is also common practice for spammers to insert randomly generated characters in the email header information. By inserting these random characters the messages are made to appear unique. This technique is used to mask the repetitive nature of the attack, and in some instances these messages will escape detection from basic spam scanning and email filtering tools. It is also believed that such strings may be used to validate address databases.

Poor User Practices

8. Another way that individuals fall prey to the spamming community is they make the mistake of opening a spam email. Many spam messages include some hidden code (often in the form of an embedded script in HTML) that reports back to the originator when the email is opened. The success of this tactic will largely

depend on the target's email client and how it is configured. Ideally the client will be configured to reduce or avoid the use of HTML email, thereby preventing the embedded scripts from reporting back to the originator. In many cases however these preventative steps are not taken, and when the email is unwittingly opened the victim's address is authenticated as valid. This address then becomes a valuable spammer commodity. By opening a spam massage three things can be accomplished:

- The victim's email address is validated as active to the spammer
- The victim is identified as a potential consumer having taken the time to open and read the message
- The victim's lack of user knowledge is demonstrated indicating they are likely to be further tricked into surrendering additional personal information that could be used to steal their identity

9. Besides these pitfalls, users also fall prey to "opt-out". Generally, there will be a tag somewhere on the spam message that tells the user that they can have their email address removed from the spammer's database by clicking on the tag. This is most often a trick used by the spammer. Clicking on the tag sends a reply to the spammer that verifies the email address is active, thereby encouraging the delivery of more spam. Spammers will often generate lists of valid email addresses through these methods, and sell them to other spammers who re-use them again and again. Some anti-spam legislation in the UK is focusing on the unsubscribe option in an effort to protect computer users from those companies that fail to honour their opt-out commitment.

10. Most mail server software can be configured to keyword search outgoing mail. By enabling this feature a system administrator may be able to block a return message to a spammer. By creating a list of keywords or phrases that are likely to be found in a reply, and assigning a threshold value to each one, a system administrator may be able to quarantine outgoing spam replies and thereby prevent users falling prey to such ruses.

Purchased Address Lists

11. Since these email address lists are perishable they are urgently marketed through numerous web sites on the Internet. Millions of addresses are available and they are typically packaged onto CDs and can be mailed directly to any customer. Since the quality of these email addresses is measured by their age the cost varies. However, they are typically in the price range of £30 to £175 for millions of addresses that are advertised between 24 hours and a month old. Once an email address has been harvested it will tend to be recirculated over and over regardless of whether it has been re-verified as active.

Automated Address Harvesters

12. Spammers will also use automated tools generally referred to as harvesters to build a target email address lists. There are many of these automated products available through the Internet. Most are designed around a couple of similar themes of data collection and navigation. The most common are referred to as "spam bots" and spiders or web crawlers. Research carried out in the USA in 2002 demonstrated that the vast majority of spam targeted addresses had been posted on public websites.

13. A spam bot is designed to enter a web page and scan for hyperlinks and email addresses. Upon completing this task, the spam bot may take different actions depending on the sophistication of the program. The more advanced programs will initiate steps whereby the newly identified email addresses will begin receiving spam immediately. Other programs that are less evolved will simply store the addresses for later use. After this retrieval process has been completed on the initial web page, the spam bot will look to the list of hyperlinks collected, proceed to one of these hyperlinked web pages, and start the process all over again.

14. Spiders/web crawlers are more general purpose programs that were originally designed to serve as an integral part of an Internet search engine. All of the common search engines used on the web today employ spiders. These programs are continually navigating through the Internet and are designed to analyze web sites and index certain topics. In principle they are identical to spam bots, as the method of navigating to web pages that are of interest to a spammer is the same. A spam bot relies on a list of hyperlinks contained in the initial target page. Since spam bots can be optimised to search for specific data sets, they are an attractive tool for spammers that are looking to build an email address repository. Spammers will seek web sites that are likely to contain large email listings such as Usenet news groups, or any public discussion forum. Therefore, upon locating the web pages with the highest probability of success, the spam bot will extract email addresses from only these sites. Spiders are more sophisticated programs than spam bots, the major difference between them being spiders are designed to visit many sites in parallel.

15. There are numerous email address extraction tools for sale on the Internet. These tools are designed to be compatible with most Microsoft Windows operating systems. Most are priced at less than £175 and are designed with a user friendly point and click GUI, giving even the most novice spammer the capability to launch an attack.

16. When the spammer has either purchased or collected a sufficient email addresses, they face the problem of distributing millions of email messages while protecting their identity as the originator. Unfortunately, this is made all too easy through the use of open relay mail servers and open web proxy servers.

Open Mail Servers

17. To send and receive email a computer must be connected to a mail server. A mail server is simply a computer that is connected to the Internet that is running software that allows it to process email. Ideally a mail server will be configured in a secure mode. When email is sent through a secured mail server, the software compares the sending address to a table of users allowed to send email (who may be members of a particular domain). If the sending address is found on the table the email is processed. Similarly, when an email is received, the mail server will again compare the destination address to a table of legitimate recipients and if found the email will be delivered. Additionally, some mail servers can be configured to require authentication prior to sending or receiving email.

18. Not all mail servers are configured as described above. Some are configured with their settings open, which means the mail server will forward any email it receives without comparison to a table of legitimate users. These mail servers are referred to as open relays. Since open relay mail servers are configured to accept and transfer email on behalf of any user, including unrelated third parties, they are used heavily by the spamming community. In so doing, they distribute their email

consuming the bandwidth of the targeted open relay, thereby increasing the response time for the legitimate user.

19. Besides ensuring a mail server is configured properly to prevent such abuses, many mail server software developers have incorporated anti-spam features into their products. Information on many of these products, that give system administrators the option to attempt spam filtering of incoming email at the server leve,l are available through http://spamcon.org/directories/server-filters.shtml

20. Public awareness is increasing regarding the pitfalls of maintaining an open relay server on the Internet. However, a simple Internet search for "open relays" will yield thousands of servers that are still operating with their settings open. As is typically the case when dealing with any underground group such as spammers, they are generally one step ahead of general public awareness. Spammers have been branching out and are using open web proxy servers that support or have access to email functionality.

Open Web Proxy Servers

21. Generally a web proxy server is nothing more than a computer that is running specific software that provides for controlled access to the world wide web from behind a firewall. From an internal network, the web proxy server listens for client requests and then forwards them along to remote web servers outside the network firewall. The responses are received again by the proxy server and relayed to the client. Most institutions use web proxy servers to improve the response times of web queries. This is accomplished because the web proxy server can store copies of previously retrieved web pages for a predetermined amount of time. Consequently, identical requests from different clients can be processed internally, thereby saving the time that would otherwise be spent going back to the web. Besides increasing performance, web proxy servers are also used as a measure of control. By engaging the content filtering features of the proxy software, system administrators can generally limit access to certain types of web sites. Additionally, a web proxy server can be configured to keep a log of a user's activities and therefore be used as a means to identify inappropriate conduct.

22. As with mail servers, web proxy servers are often configured poorly, thereby offering services to a wider audience than is intended, and it is here that the spammer takes advantage. The problem with open web proxies is that anyone on the Internet can use them as go-betweens to perform just about any action related to web access.

23. Open source research indicates that spammers are using open web proxies at an increasing rate. This may be attributed to increased awareness regarding the pitfalls of maintaining an open mail server. Once a spammer has access to an open web proxy, they can do the following activities with little risk of being traced back beyond the network address of the open proxy:

- Send mail from unsecured form-to-mail scripts on that or any other web server
- Connect to thousands of throwaway "freemail" (Hotmail, Yahoo, etc.) web mail accounts per minute and send potentially millions of spam messages

24. Depending on the sophistication of the proxy it may also be possible to send a mail message to a mail server via a web proxy. The web protocol, HTTP, is a plain text protocol, just like the email protocol, SMTP. A web proxy that does not check the application content may act as direct spam relay.

Other Proxies

25. It is not only mail and web proxies that are open to abuse. It is also possible to misuse open proxy servers for other types of Internet traffic, an example being a SOCKS proxy server which in the past was often used by firewalls to translate addresses and ports to hide details of the internal network from the Internet. Just like a web proxy, if the proxy does not check the application content, it may be used as spam relay.

Malicious Code

26. Open source reporting indicates that spammers are adopting the use of malware to create open proxies that can be used to distribute their messages. Increased public awareness of the many negative aspects of spam has lead many organisations to apply anti-spam technology to their networks. Consequently, some avenues through which spammers distribute their messages have been shrinking. In reaction to these events, spammers have sought other methods to create opportunities to continue their message distribution.

27. There have been several new trojan horse programs detected that are designed to specifically aid the spammer. Symantec recently reported on a trojan named Backdoor.Migmaf that is a reverse proxy trojan horse that redirects HTTP requests to a master web server. This enables the creator of the trojan to hide the real IP address of the web server. Analysis performed by the Managed Business Solutions Corporation LURHQ Threat Intelligence Group indicates that anyone requesting the URL core.onlycoredomains.com (among several others) would be directed to one of the trojaned machines, where the connection would in turn be relayed to the master web server. The returned page would be passed back to the trojaned machine, where it was then sent back to the requesting user's web browser. In this way it is impossible for the user to tell the actual IP address of the master server, thereby giving the spammer's real web site refuge from being shut down by their ISP.

28. Migmaf also listens on TCP port 81 and acts as a SOCKS proxy server. This allows the spammer to send spam email through the trojaned machine to any number of target recipients. This means the spammer has a complete end-to-end anonymous system for spam. Migmaf reportedly has no spreading capability. It has been suggested in open source reporting that spam friendly trojans such as Migmaf could be distributed as a virus payload. There are many other examples of this type of code that are being used to open avenues for spammers to operate. It is likely this technique will be used increasingly as spammers attempt to stay one jump ahead of the current mitigation techniques.

Hijacking Registered but Unused Network Addresses

29. The Border Gateway Protocol (BGP) is protocol for routing network traffic between autonomous systems. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is a protocol used between ISPs to announce and withdraw routes.

30. An increasingly common spamming technique is to hijack network addresses that have been registered but are not in use. Spam email can then be sent from these addresses. The hijacking is achieved by announcing a route for the unused network addresses and withdrawing the route when the spam has been sent (which

may be a matter of minutes or seconds). See http://www.nanog.org/mtg-0310/pdf/hutzler.pdf for more details

31. To reduce the risk of this type of hijacking, ISPs should filter announcements received from, and sent to, BGP peers, in much the same way that local ingress and egress filters are applied to individual networks. ISPs may wish to monitor announcements of unused network address ranges and complaints of spam relating to those address ranges. (The originating autonomous system can be traced by looking at the autonomous system path in the BGP routing table.)

Anti-SPAM Techniques

32. Besides enabling the security features included with most mail and web proxy server software, there are a vast number of anti-spam products and subscription services available to aid system administrators. These have been developed as a result of increased public awareness regarding the threat that spam presents to business productivity and security. Most of the products available today are centred around a few common techniques that will be covered in detail in the following sections. They include:

- Blacklisting (negative filtering)
- Whitelisting (positive filtering)
- Greylisting
- Bayesian filtering
- Heuristic or rule based filtering
- Sender Policy Framework (SPF)

Blacklisting

33. A blacklist is a database that can include IP addresses, domain names, email addresses, or email header content. Although blacklists can be configured to use any combination of these, they are most commonly composed of open relay mail server and open proxy server IP addresses. The following discussion focuses on IP address-based blacklists. These open servers are frequently used to send spam. Blacklists are generally provided to the public by a third party as a subscriber service, and can be used through many mail server software packages to help prevent the delivery of unwanted mail. The majority of organisations that use this method of negative filtering will use a third party provider. The primary reason is that, as spammers' intrusion techniques evolve as fast as the prevention techniques, maintaining a current blacklist is a full time job. Generally mail server software is configured (e.g. Sendmail's rulesets) to give a system administrator the option to consult a blacklist automatically. Most of the subscription services continually update their databases as the spamming community is constantly changing. Those lists that are continually updated are also referred to as Real-time Blackhole Lists (RBLs). Using an RBL is most commonly accomplished through two different methods, transfer or inquiry.

34. In inquiry mode the subscriber must know the IP address of the originating host or mail relay, and then use some network protocol to query the blacklist service provider to determine whether the host is listed in the RBL. In transfer mode on the other hand, a copy of the entire RBL is downloaded to a host that is owned by the subscriber. Many of the blacklist subscriber services provide automatic updates to the RBL. Therefore, queries can be made locally within an organisation.

35. Regardless of the method, the sending IP address is compared to a blacklist. If there is a positive result, the subscriber can choose any action which is appropriate to their site security policy. Generally, the two options are to bounce the mail or place it in a quarantine folder that will be periodically checked to retrieve any messages that were improperly graded. This type of filtering is often used by various ISP and bandwidth providers to filter out spam sent across their networks or to their subscribers. System administrators who use this type of service should report problematic domains, autonomous systems and addresses to their service provider so they may be added to the blacklist.

Anti-Spam Sites

36. The following list is composed of anti-spam sites. Many have links to search their blacklist databases.

- SpamCop blacklist http://spamcop.net/bl.shtml
- MAPS blacklist http://www.mail-abuse.org/cgi-bin/lookup
- Open Relay blacklist http://www.ordb.org
- Relay Stop List blacklist http://relays.visi.com/
- Distributed Server Boycott List blacklist http://dsbl.org
- WireHub blacklist http://basic.wirehub.nl/spamstats.html &
- http://informatie.wirehub.net/error/dynablock.txt
- SPAMHaus blacklist http://www.spamhaus.org
- SPEWS.org blacklist http://www.spews.org/
- Osirusoft blacklist http://relays.osirusoft.com/cgi-bin/rbcheck.cgi?addr=
- Monkeys.com blacklist http://www.monkeys.com/anti-spam/filtering/
- Blitzed.org blacklist http://www.blitzed.org/ Has to do with IRC (chat)
- WebTV blacklist http://info.webtv.net/spam/
- SPAMbag blacklist http://www.spambag.org/query.html

37. There are numerous Internet organisations, perhaps most prominently the Mail Abuse Prevention System (MAPS) (http://mail-abuse.org/rbl/), who maintain lists of individual IP addresses and various netblocks that are known in some way to support spammers, having open relays, hosting web sites, distributing marketing spamming software, etc. All MAPS lists are based on IP addresses only, not domain names, email addresses, URLs, or message contents. The MAPS web site offers the instructions and syntax required to allow a user with good computer skills to integrate Sendmail with the MAPS RBL. This will benefit system administrators that are hosting mail servers using UNIX based operating systems. (Sendmail is a popular UNIX-based email server software.)

Drawbacks to using Blacklists

38. Blacklists are highly controversial because in some instances they prevent legitimate email from being delivered. The following are some of the most common failings of the blacklist approach to spam filtering.

- Spammers are always changing their addresses. By using “throwaway” dialup accounts, DSL or cable modem addresses, many spammers stay one step ahead of the blacklist updates. Even if the blacklists are just a day behind, the spammers can get their messages out.
- Suppose some user@ISPx.com starts sending spam, and a spam recipient complains to a blacklist service provider. Ideally, just that single address would be added to the blacklist. However, the entire domain will sometimes get added to the blacklist. Thus email from allusers@ISPx.com is labelled as spam and legitimate email will not be delivered.

- Someone might falsely report user@ISPx.com as being a spammer. In the world of blacklists, you are guilty until you prove yourself innocent. If a user is falsely accused they can be stuck on a blacklist for days while the error is corrected.
- One of the most notable problems with blacklists is they are used by ISPs to filter out spam. Legitimate email can be blocked by the ISP without the intended recipient ever knowing that the email was sent.
- Blacklisting can sometimes be defeated by spammers who forge addresses that belong to legitimate correspondents
- Blacklist service providers have been targeted with DDoS attacks initiated by spammers in retaliation for limiting their ability to deliver their junk mail. These attacks will limit or in some instances prevent the service provider from maintaining contact with their customers. The industry consensus is these types of attacks are likely to increase in the near term.

Whitelisting

39. Whitelisting is a very simple approach to spam filtering. A whitelist is composed of all the email addresses that any particular organisation has determined are valid. Only email that is received from one of these valid addresses will be accepted. Generally a whitelist would consist of every email address in all users address books, contact lists and corporate directories. Whitelists are often automatically supplemented with recipient addresses of outgoing mail messages that may not appear in a user's address book. If an email is received that is not on the whitelist then it can be bounced or placed into a quarantine folder to be reviewed on the prospect that some legitimate email was improperly handled. Another whitelist strategy is to take messages from unknown senders and hold them in a pending queue until the sender responds with a confirmation or some additional information that validates them as a legitimate sender. Typically spammers operate under a fire and forget mentality and will not take the time or have the resources to respond to such a challenge.

Drawbacks to Whitelisting

40. Whitelists are not completely reliable because in many instances they prevent legitimate email from being delivered. The following are some of the most common failings of the whitelist approach to spam filtering.

- Unlike blacklists, whitelists are created and maintained by individual organisations, and then require a reasonable amount of maintenance on the part of that organisation's system administrator to keep them current
- An organisation is likely to receive many legitimate emails from people who are not whitelisted, so a pure whitelist spam filter will generate a large number of false positives
- Whitelisting places the burden of managing a quarantine database on the receiving organisation. Manual checking is a time consuming process that costs organisations money because it is labour intensive
- Whitelisting can be defeated by spammers who forge addresses that belong to regular correspondents
- The performance requirements of a positive filtering implementation can be fairly significant. A typical user may contribute several hundred addresses to the organisation's whitelist after using the system for a prolonged period. The processing requirements needed to filter each incoming mail message against an organisation's whitelist can be considerable, resulting in the need for additional mail server hardware.

Greylisting

41. Greylisting is a cross between black and whitelisting. The leading developer of the greylisting techniques is Evan Harris (see http://greylisting.org/). A key element of the greylisting method is its automatic maintenance features. The Greylisting method relies on three pieces of information about any particular mail message which are referred to as a triplet. They are:

- The IP address of the host attempting the delivery
- The envelope sender address
- The envelope recipient address

42. This unique triplet is used to identify all email messages and it is applied to one very basic rule which is:

- If this triplet has never been seen then refuse the delivery with a temporary failure, in addition to any others that may come within a certain period of time

43. The email protocol SMTP is considered an unreliable transport. Therefore, the possibility of a temporary failure is built into the core specifications (for more information see Request For Comments (RFC) 2821 at http://www.ietf.org/rfcs/rfc2821.html). As such, any well behaved mail server should attempt a series of retries if given an appropriate temporary failure code for a delivery attempt. According to the testing performed by the greylisting developers in mid 2003, the vast majority of spam appears to be sent from applications designed specifically for spamming. These applications appear to adopt the fire and forget methodology. That is, they attempt to send the spam to one or several Mail Exchange (MX) hosts for a domain, but then never attempt a true retry as a real mail server would. Developer testing revealed over 95 percent of test spam messages sent were blocked, with no legitimate mail ever being permanently blocked.

44. There is some evidence to suggest that spammers may be using viruses to deliver trojan code that will then allow a spammer to convert a compromised host into an open proxy. Thereby creating an avenue to deliver more spam while hiding the identity of the originator. Greylist developer testing indicates this technique to be extremely effective in blocking these types of malware, as they generally do not tend to retry deliveries. Since these viruses are fairly large, bandwidth and processing savings are significant versus the standard method of accepting delivery and local virus/malware scanning.

45. This approach comes with a minimal price in terms of resources when considering the use of a local data store for the triplet and other metadata. However, there is no required network traffic caused by greylisting other than that which is associated with the connection itself. This is because the contents of the message are not checked, so there is very little processing overhead unlike many other spam blocking methods.

46. Since the greylisting technique delays acceptance of unknown mail it does generate more work for the mail server that sends a legitimate email. However, it will also generate more work for the spammer's systems should they elect to make another delivery attempt. If enough organisations adopt the greylisting technique, then spammers would increasingly have to retransmit their messages thereby placing an additional resource drain on their operation. It has been suggested in open source reporting that if enough system administrators adopt this technique it may raise the threshold of difficulty to deliver spam beyond the limit of some

spammers. This may be true: however there are too many variables with unknown values (all of the spammers associated costs of running a spam operation, spammers success rate in selling their product or service, price of the product or service, profit margin, etc.) to substantiate this claim.

47. Because this is a relatively new technique and not widely used it is unclear how sustainable it will be over time. The spamming community has proven their ability to adapt to constraints over time. This technique does however show promise as one of several methods that could be used in unison because its thresholds can be modified to counter the evolving efforts of the spamming community. For more on the full greylist implementation see http://greylisting.org.

Bayesian Filtering

48. Bayesian filtering is a relatively new technique that is receiving a lot of attention from anti-spam software developers. The recognized pioneer in this new field is Paul Graham who has developed and applied a statistical analysis model (See: "A Plan for Spam" http://www.paulgraham.com/spam.html) to grade each piece of email using a series of algorithms. The probability of an email being spam is calculated based on the message content including the entire text, headers, embedded HTML, and JavaScript. Alphanumeric characters, dashes, apostrophes, and dollar signs are considered to be part of tokens, and everything else to is taken to be a token separator. Tokens that are all digits and HTML comments are ignored. Because this method uses probabilities, it considers all the evidence in the email, both good and bad. Words that occur rarely in spam like "though" or "tonight" or "apparently" contribute as much to decreasing the probability as bad words like "unsubscribe" and "opt-in" do to increasing it. So an otherwise innocent email that happens to include the word "sex" is not going to get tagged as spam.

49. The approach is to filter each user's email based on the spam and non-spam mail that they receive. Therefore, the spam filter characteristics will differ between users as each have their own unique behaviour patterns that apply to email. The Bayesian filter is designed to operate with increasing efficiency over time, improving as it is exposed to more and more good email and spam. To accomplish this each user should have two delete buttons, "delete as good email" and "delete as spam". Anything deleted as spam goes into the spam profile, and everything else goes into the non-spam profile. This lets each user decide their own precise definition of spam, and makes it harder for spammers to tune emails to get through these many individually tailored filters.

50. A clear advantage to this filtering approach is that it is dynamic, evolving with the changing characteristics of spam. For example, should spammers start using "$ex" instead of "sex" to evade detection, the filter would become attuned to this change. Therefore, the word "$ex" would be regarded as far more compelling evidence to condemn the message as spam than if a message simply contained the word in its unaltered form.

51. One of the next steps towards refining the Bayesian technique is to filter based on word pairs or word triplets, as opposed to the current model of single word filtering. Word combinations are expected to yield much clearer estimates of probability. For example, the word "offers" reportedly has a spam probability of .96. If word pairs are considered as a basis to determine probabilities such as "special offers" and "valuable offers" the probability increases to .99. These refinements and others such as weighted focus on specific parts of the email message are planned for the next revision of the Bayesian filtering model.

52. Though this approach to spam filtering does show promise there is some evidence to suggest the spamming community is adapting. Spammers will sometimes include neutral words in their messages in an attempt to minimise the benefits of this technique. Therefore it may serve best as one component of a multi-layered approach to spam filtering.

Heuristic (Rule Based) Filtering

53. Heuristic filters look for patterns that indicate spam such as specific words, phrases, lots of uppercase characters, exclamation points, unusually configured headers, dates in the future or the past, etc. The performance of heuristic filters varies widely. The most basic simply filter on word sets in the message body, and are typically easily beaten by spammers. These unsophisticated filters also tend to have a high false positive rate. However, well developed rule based filters like SpamAssassin can be quite effective. SpamAssassin is designed to interface with other spam filtering techniques as well. SpamAssassin is an open source product designed for UNIX operating systems, although there are commercial versions as well. The SpamAssassin home page (http://spamassassin.org) has links to downloaded UNIX-oriented front-end scripts, versions for Windows, as well as commercial versions. The spam identification tactics used in SpamAssassin include:

- Header analysis: Spammers use a number of methods to mask their identity in an attempt to overcome spam filtering. SpamAssassin examines header data, filtering against known spammer profiles
- Text analysis: SpamAssassin looks at the message text, filtering against a known spammer jargon
- Blacklists: SpamAssassin supports many existing blacklists
- Vipul's Razor: SpamAssassin is configured to integrate with Vipul's Razor (http://razor.sourceforge.net/) which is a distributed, collaborative, spam detection and filtering network. Through user contribution, Razor establishes a distributed and constantly updating catalogue of spam in propagation that is consulted by email clients to filter out known spam. Detection is done with statistical and randomized signatures that efficiently spot mutating spam content. User input is validated through reputation assignments based on consensus on report and revoke assertions which in turn are used for computing confidence values associated with individual signatures.

54. The advantage of rule-based filters over methods such as Bayesian filters is that they are easy to install at the mail server level. Bayesian filters for example require users to train them by telling them when they misclassify an email, so running one on the server is more complicated.

55. The main disadvantage of rule-based filters is that they tend to have high false positive rates. Many of the attributes of legitimate email are cleverly incorporated into spam mail. Therefore, when filtering it is inevitable that some legitimate email will be tagged as spam. One way spammers accomplish this is to add random words to HTML email using white text so the victim will have no visual indicator that the message has been tailored to beat a spam filter. These words are randomly generated and therefore change with each new message, thereby increasing the chance of success for the spammer. Rule-based filters such as SpamAssassin can be configured to trap for white text. However, as described above there will be instances where legitimate HTML email will contain white text on coloured background and therefore produce a false positive.

56. Another disadvantage is that the rules are static. When spammers modify their tactics, the software developers have to write new rules to catch them. Because

rule based filters are static targets, spammers can tune their emails to get past them. Sophisticated spammers already test their messages on popular rule based filters before sending them. In fact, there are sites that will do this for free. Therefore this technique will work best for organisations using third party service, possibly augmented by local rulesets based on organisation-specific problem mail.

Sender Policy Framework (SPF)

57. SPF is not a commercial product; it is an open standard, intended to be an extension to SMTP that will prevent forgery or spoofing of the sender's email address (a technique commonly used by spammers). An Internet draft has been written (see http://spf.pobox.com/spf-draft-20040209.txt). SPF is not a spam-filtering technology, although its functionality will certainly help reduce the amount of spam traversing the Internet. Instead of analysing the content of messages to identify spam, SPF allows Internet domain administrators to describe their email servers in an SPF record that is attached to the Domain Name System (DNS) record. Other Internet domains can then reject any messages that claim to come from that domain but were not sent from an approved server.

58. Unlike spam filters, the SPF technology allows email gateways to analyse the email envelope, a wrapper for the message that is transferred between mail servers before the full message is sent. Messages that do not come from a valid server at the domain can be dropped, before any message content is sent. Because no message content is sent, organisations save network bandwidth and computing resources compared with filtering, which requires bogus messages to be sent, received and then analysed.

59. SPF is designed to tell the user the following:

- The sender is good the sender has previously announced that they do send mail from that IP address
- The sender is bad the purported sender has published a list of IP addresses they send mail from, and the client IP is not one of them
- The sender may be good or bad. The sender domain is in a transitional phase; it is methodically converting its users to be SPF compliant, so the community should go easy on any violations for the present
- SPF doesn't know. The sender has not published any IP addresses, so the message could be legitimate or not

60. Like so many of the filtering techniques, SPF will require widespread use for it to have a meaningful effect on the spamming community. Gaining the cooperation of all ISPs will be challenging. Those legitimate service providers that do not subscribe to SPF will serve as avenues through which spammers will still deliver their unwanted junk email.

User Recommendations and Countermeasures

61. The following list of recommendations should be followed by all organisations. However it is geared towards those organisations that have a limited budget.

- Delete spam messages. If possible delete spam messages without opening them. Think of it in the same context as throwing away the junk mail received at home.
- Never reply to a spam message. This can inadvertently cause all original addressees to receive the reply causing another flood. This result is true most often with chain letters than with run of the mill spam.

- Never respond to “instructions to remove me from the mailing list”. Most often the spam victim will receive a bounced mail message in reply. Typically, taking this action will add a victim’s address to many spam lists, as it serves as a confirmation that the account is active and the email is being read. There are legislative efforts underway to force organisations to adhere to their opt out commitment.
- Best practice is to never post to a news group or bulletin board. However, if compelled to do so, post messages using a modified email address that will never yield a reply through automated means. For example, the following legitimate address “user@isp.com” could be easily modified to “user((at))i-s-p-dot-c/o/m”. This will at least give the user a fighting chance at avoiding detection by automated harvesters.
- If hosting a webpage or web site, do not post any legitimate email address as a hyperlink. It’s a bit more work, but use a form instead to act as an intermediary so that a legitimate email address is never revealed.
- Users should remove or modify their signature block when sending email or posting to any open forum. This sounds very basic but often people will forget that this automatic signing feature is populated with precisely what the spammer is looking for.
- Consider using a throwaway email account and use that address only when posting messages. If a throwaway account is used, be prepared to dump it for a new one when that inbox begins filling with spam.
- Throw away addresses can be tailored for specific events. Consider a user that will attend a business meeting and knows they will be giving their email address to new associates. Prior to attending the meeting the user creates a tailored address named after the meeting itself “meetingJanuary2004@isp.com” and distributes it only to the members of this business. This technique can be used to communicate with new business associates until they have established their credibility. Should the user receive spam at this disposable address it will point directly to the group responsible.
- If operating a mail server, ensure that it is configured to allow only legitimate clients to send and receive email.
- If operating a web proxy server ensure that it is configured to prevent unintended uses.
- System administrators should consider enabling the anti-spam features that are incorporated into many of the server software packages currently on the market
- Be selective about how email addresses are distributed, and encourage individual end users to be similarly selective
- Never list all the email addresses of an organisation on a single web page, and where possible do not list them as direct hyperlinks

System Administrator Recommendations

62. There are a wide variety of anti-spam techniques available. Selecting the most appropriate method will depend on a variety of factors such as: type of network architecture, organisation size, level of user knowledge and proficiency, and budget.

63. Assess the skill level of an organisation’s average user and institute spam filtering policies accordingly. Organisations with a high skill level could expect individual users to manage their own spam filtering regime. Organisations that have employees with a lower skill level should consider spam filtering at the server level. In either case the goal of every organisation should be to increase user awareness through training.

64. Small organisations that have a limited amount of email traffic and limited system administrator support should consider simple whitelisting. It is an easily installed and a relatively inexpensive method that will eliminate a major percentage of unwanted email. Its main drawback is a high false positive rate requiring a manual check of a quarantine folder.

65. Large organisations that communicate heavily through email and that have a large and well funded system administrator team should consider integrating different techniques. A combination of whitelisting and greylisting or Bayesian filtering as an element of heuristic filtering will most likely yield good results. The drawback to these approaches is that they cost more to configure and maintain, and require system administrator staff with deep technical knowledge.

Considerations While Shopping For Anti-Spam Products and Services:

66. The following are considerations when choosing an anti-spam product:

- What is the claimed spam detection rate?
- Have claims of product effectiveness been verified by a reputable third-party?
- Which techniques (blacklisting, whitelisting, greylisting, Bayesian filtering, heuristic filtering) are supported?
- Is it possible to modify the system configuration to improve spam detection?
- If so, what level of expertise is required to maintain and modify the configuration?
- If the anti-spam tool is managed by a vendor what are the minimum and maximum response times for a service request such as changing the configuration?
- Is the product easy to use?
- Are there any minimum service guarantees that come with the product?
- Is the product optimised for the customer network operating system?
- Are there any advantages in using an open source product considering the level of system administrator expertise within the customer organisation?
- What level of customer support does the product include, for example is there a 24 hour help line?
- Do the contract terms fit the customer organisation business model and are there price reductions for extended contracts?
- What are the penalties for breach of contract should the product or service be inadequate?
- What is the cost of the product or service and what are the projected costs of installation and maintenance?
- Is the product reliable?
- What are the bandwidth requirements for the product or service?
- What additional load does the application place on existing equipment?
- Will the product or service require the purchase and installation of any additional computer hardware?
- Will end users require training to use the product and what are the associated costs?

References:

1. Federal Trade Commission - Facts for Business. Open Relays - Close the Door on Spam. http://www.ftc.gov/bcp/conline/pubs/buspubs/openrelay.htm
2. Lencom Software Inc. Fast Email Extractor 4.4. http://www.email-marketing-easy.com/FEE.html

3. Mercator: A Scalable, Extensible Web Crawler by Allan Heydon and Marc Najork, Compaq Systems Research Centre.http://research.compaq.com/SRC/mercator/papers/www/paper.html
4. Tconsult Inc. What are Spiders? http://www.tconsult.com/faq/spider.aspx
5. SpamBot Beware, Background and information. http://www.turnstep.com/Spambot/info.html
6. SPAMHAUS - The Spamhaus Project. http://www.spamhaus.org/news.lasso?article=6
7. What is a Dictionary Attack. http://www.filterpoint.com/help/dictionary.html
8. Fast MLM Leads.com. http://www.fastmlmleads.com/email.html
9. OneBIGworld.com. http://onebigworld.com/usadirectory/metasearch.cgi
10. mail Unknown anti-spam software. http://www.mailunknown.com/Opt-Out.asp
11. TidBITS#704/03-Nov-03. http://www.tidbits.com/tb-issues/TidBITS-704.html
12. Purveyor administrator's guide - web proxy servers. http://vms.process.com/~help/helpcgi.html
13. Email Blacklist Directory. http://www.spam-blockers.com/SPAM-blacklists.htm#what-is-a-blacklist
14. http://www.spamcon.org/directories/server-filters.shtml
15. http://www.spam-blockers.com/SPAM-blacklists.htm#what-is-a-blacklist
16. MAPS blacklist http://www.mail-abuse.org/cgi-bin/lookup
17. The Electric Editors - Slaying Spamhttp://www.electriceditors.net/faq/spam.htm
18. Whitelisting Could be Spam readyhttp://www.nwfusion.com/columnists/2002/1209kobielus.html #
19. User Whitelist Spam Filter. http://www.sambar.com/syshelp/whitelst.htm
20. The Next Step in the Spam Control War: Greylisting, Evan Harris. http://projects.puremagic.com/greylisting/
21. Symantec, Backdoor.Migmaf

    Http://www.symantec.com/avcenter/venc/data/backdoor.migmaf.html

22. Rise of the Spam Zombies, Kevin Poulsen, SecurityFocus http://www.theregister.co.uk/content/55/30414.html
23. Trojan turns victims into DDoS, spam zombies, John Leyden http://theregister.com/content/56/31801.html
24. Sobig.F Illustrates Anti-Spam, Anti-Virus Convergence. by Scott Bekker http://www.entmag.com/news/article.asp?EditorialsID=5930
25. LURHQ Threat Intelligence Group. http://www.lurhq.com/migmaf.html
26. Spam Filtering-How to Use Heuristic Filtering http://files.altn.com/HowTo/SpamFiltering-HowtoUseHeuristicFiltering.html
27. Paul Graham, Stopping Spam http://www.paulgraham.com/stopspam.html
28. Paul Graham, A Plan For Spam http://www.paulgraham.com/spam.html
29. SpamAssassin Home Page http://spamassassin.rediris.es/index.html
30. http://www.cs.tut.fi/~jkorpela/html/iframe.html
31. http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm
32. http://spf.pobox.com/
33. http://www.geocities.com/spamresources/filter-bl.htm
34. http://www.nanog.org/mtg-0310/pdf/hutzler.pdf.