Global Information Assurance Certification Paper

Danny Lynn
Version 1.4

Spam-Spam-Spam

Abstract:

Unsolicited automated email, Spam, is a growing Internet threat. Though almost everyone agrees that something should be done about it, the solution to the problem remains as elusive as ever. Systems administrators have a responsibility to reduce Spam in the workplace. The cost of fighting Spam far exceeds the cost to produce Spam. Many organizations are focused on the control of Spam either through technical or legislative means. Governments throughout the world have passed various forms of anti-Spam legislation, but with no effect on overall Spam volume. Technical solutions appear the most promising in the control of Spam. There are two primary technical methods of Spam control: Blacklists and Filters. While Blacklists are the most widely utilized form of Spam control, intelligent filters are starting to arrive on the market that appear quite effective in the discerning Spam messages from the users legitimate email. There are specific measures that systems administrators can follow to mitigate the problem at their workplace, but completely eliminating Spam is an unrealistic goal.

What is Spam?

Spam is difficult to define. One person's Spam may be another's diversion, and it is this dilemma that makes it so difficult to legislate and mitigate. It seems that everyone knows what Spam is when they see it, but when asked to define "Spam", it becomes clear that almost every definition relies on the subjective view of what is not Spam rather than what is Spam. Those who understand logic can appreciate how difficult it is to define something by what it is not. So, to say that Spam is "unwanted, unsolicited email" is to say very little about what Spam is.

According to the Mail Abuse Prevention System (MAPS), an email is Spam "IF: (1) the recipient's personal identity and context are irrelevant because the message is equally applicable to many other potential recipients; AND (2) the recipient has not verifiably granted deliberate, explicit, and still-revocable permission for it to be sent; AND (3) the transmission and reception of the message appears to the recipient to give a disproportionate benefit to the sender." [1] What does this definition mean? The three clauses of the MAPS definition appear written more to protect the rights of the recipient than to provide a definition of Spam. The result of this is that the MAPS definition does not clearly delineate between Internet email Spam and other forms of communications harassment. This definition can be applied to almost any form of unsolicited or unwanted communication. It is equally applicable to Usenet postings, IRC, telephone marketing, US Postal Service bulk mail and junk faxes.

Paul Graham, a noted anti-Spam author and programmer, proposes that Spam simply be defined as “unsolicited automated email”. He writes that the salient aspect of Spam is that it is automated, not that it is unsolicited. “To start with, spam is not unsolicited commercial email. If someone in my neighborhood heard that I was looking for an old Raleigh three-speed in good condition, and sent me an email offering to sell me one, I'd be delighted, and yet this email would be both commercial and unsolicited.” [2] This example presents a much broader, more practical definition than the MAPS definition. His definition would include email sent by companies that have an existing relationship with the recipient. Graham asserts that “buying something from a company, for example, does not imply that you have solicited ongoing email from them. If I order something from an online store, and they then send me a stream of spam, it's still spam.”[3] More to the point is that Graham’s definition looks at the intent or context of the message rather than merely at the delivery mechanism. This highlights a fundamental distinction between two Spam mitigation approaches – filtering versus blocking.

For the purposes of this paper, the term Spam will refer to messages transported via the Internet using the SMTP protocol. It can be called unsolicited bulk mail (UBM), excessive multi-posting (EMP), unsolicited commercial email (UCE), or any other form of unsolicited automated email.

Where did the term “Spam” come from?

The term “Spam”, as it relates to the Internet (not the canned meat product), got its start in the timeframe of the late 1980s to early 1990s with the Multi-User Dungeons (MUD) and Multi-User Shared Hallucinations (MUSH) user groups. Apparently a MUSH user programmed a macro key to type “spam, spam, spam” in a MUSH session until the systems administrator terminated his connection. Because of his effusive, repetitive and unwanted typing, he became known as the person who “spammed” them.

The first use of Internet email as a vehicle for Spam was on May 1, 1978. A Digital Equipment Corp (DEC) sales representative wanted to advertise an upcoming sales event and attempted to send an email to all Arpanet users on the US West Coast. This misuse of Arpanet was not well received. Sixteen years later a law firm from Phoenix, Arizona posted a message to several thousand newsgroups that promoted the use of their services. “This was probably the first automated large scale commercial use of Spam, and was the incident that popularised the term, which up until then had been exclusively part of the arcane vocabulary of Multi-User Dungeons.” [4]

What is the scope of the problem?

Spam is a growing threat. According to Brightmail, Spam increased by 179% in one year, from 1.97 million messages in November 2001 to 5.5 million messages in November 2002. “This increase is even more pronounced when one considers that as recently as June 2001 the network logged fewer than 1 million Spam messages. That's more than a five-fold increase in less than 18 months.” [5] The following graph by Carl Weinschenk is from his article “Prevention is the Best Medicine: Creating an Anti-Spam

Strategy". The chart measures "Attacks" (instead of the more accurate term "Spam messages") per month.

Unique Spam Attacks June 2001 to November 2002[6]
(as Measured by Brightmail's Probe Network)

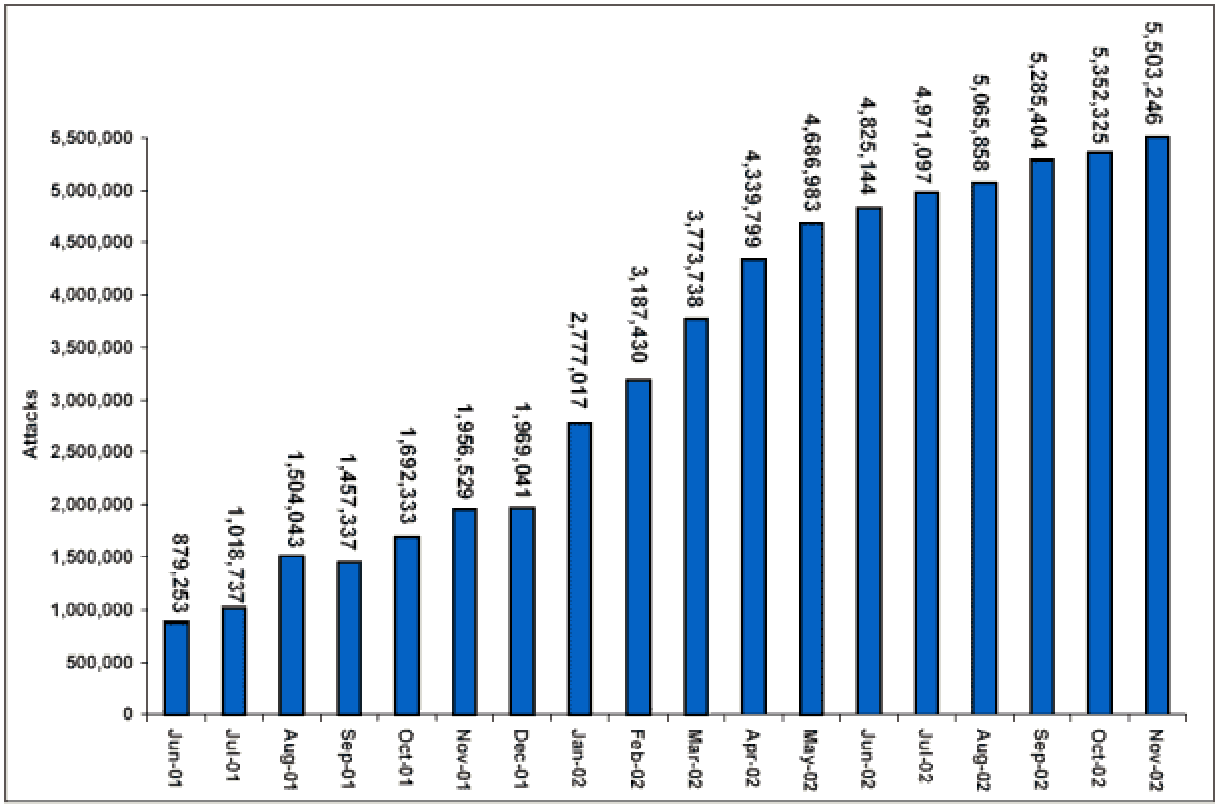

The increased volume of Spam messages over such a short period of time can be equated to a slow DOS (denial of service) attack. Unlike DOS, which is characterized by fast but small TCP communications, Spam DOS attacks characterized by slow but large TCP communications, Spam messages involve more data and therefore utilize more bandwidth and network resources.

What is the Spam economy?

As long as people purchase the products offered in Spam messages, then Spam vendors and those that request this form of marketing make money. The cost of sending Spam messages is disproportionately low when compared to the cost incurred by the Spam recipient trying to manage the Spam message. For this reason Spam is considered a "parasitic economy". In an article published by the ePrivacy Group, Spam vendors can make profits on rate of return as low as 0.001%. They cited a case in which a Spam email marketing campaign of 3.5 million messages resulted in 81 sales within the first week. This equates to a return rate of 0.0023%. Since each sale was

worth $19 to the Spam vendor, they made a little over $1,500 with few incurred expenses. [7]

The Spam message must employ a certain amount of social engineering. To be effective someone must be drawn to open the message and read its contents. The email subject must be titillating, timely, familiar, or in some way enticing. Of the various categories that Brightmail's Probe Network™ (a Spam filtering service) tracked over the 20 days surrounding Valentine's Day, the following is a breakdown of subjects found in over one million of the messages:

- o Flowers & Chocolates: 50%
- o Sexual performance enhancers: 20%
- o Lingerie: 10%
- o Dating/matchmaking services: 10%
- o Other Gifts (stuffed animals, jewelry, etc.): 10% [8]

Holidays, particularly Valentine's Day, present an excellent marketing opportunity for spammers. Virus writers and spammers can easily hide their payload within cleverly worded messages. The victims are tricked into compromising their own systems (and often those of others) on a day when it is traditional to receive unexpected messages and gifts. "Viruses such as 'Kornikova' and 'I Love You' are classic examples of how the perpetration of viruses can be dramatically increased by exploiting common weaknesses. Before thinking about the content of the email, recipients opened it only to discover that it contained a malicious payload. This then used the computer to email everyone in the address book and spread further." [9]

Why is Spam a security issue?

Spam presents one of the most difficult security problems to solve. It has an ambiguous signature and is decentralized, pervasive and persistent. Spam can be viewed, as a form of harassment and it is more than a nuisance. Spam messages often carry unwanted payloads, such as pornography, that can easily offend the recipient. Further, Spam via email is the transport method of choice for the delivery and dissemination of malicious software – viruses, Trojans and spy-ware. Simply put, an Internet threat that directly interacts with a company's employees is a danger to that company. It is the responsibility of that company's systems group to manage this threat and user expectations are high. Vince Tuesday's article titled "Spam Issue Viewed As IT Failure" points out that "the problem of unwanted e-mails may not be the biggest security threat my company faces, but it's a highly visible failure to our employees and threatens to taint the IT security group's good reputation… One of the ways that staffers judge good security is if they can work free of fear... Life is good… Apart from spam." [10]

A workplace free from harassment is a legal issue and is one way that employees measure the quality of their work environment. Employers are increasingly aware that Spam can be viewed as a form of harassment, and inaction on the company's part could be construed, at the very least, as insensitive. To underscore this point, Sharon

Gaudin wrote in a recent web article that "more pornographic spam is hitting inboxes -- many with increasingly graphic subject lines and often graphic images enclosed. It's a trend that is offensive to many users. And many IT and human resource leaders fear that the increase in pornographic email messages is creating a hostile work environment, leading many companies down a slippery legal road." [11] Many companies are implementing anti-Spam measures in response to these human resource issues. Is Spam a security threat? The answer is yes.

How to control Spam?

Legislation:

Since 1999 twenty anti-Spam legislative proposals have been submitted to the US Congress, and to date, none have been enacted.[12] Within the United States, twenty-eight states have passed anti-Spam laws.[13] In Europe sixteen countries plus the European Union have passed various anti-Spam measures.[14] Ten nations from the rest of the world have also implemented anti-Spam legislation.[15] Undaunted, Spam continues to move freely across the Internet and with increasing regularity. Given the borderless nature of the Internet and the difficulties of tracking rogue spammers through the maize of open-relays and dial-up accounts, legislation may prove to be an ineffective mechanism for the control of Spam.

Technical:

There are two basic approaches to isolating Spam from a user's email in-box. One approach is to block Spam at its source by rejecting all SMTP traffic originating from known Spam sites. The second approach is to employ message content filters. The success of either approach can be measured using a common standard. In a web article titled "Filters vs. Blacklists", Paul Graham wrote that "the real test of any technique for eliminating Spam is not how much Spam you can stop, but how much Spam you can stop without stopping a significant amount of legitimate email. That is, how do you design a defense against Spam so that the error in the system is nearly all in the direction of false negatives rather than false positives?" [16] A Spam control method that blocks 99% of all incoming Spam sounds pretty good, but if that same method also blocks 10% of all legitimate email traffic then it is not so useful.

Black and White Lists:

To date the most common approach to Spam control is to utilize a list of IP addresses that are "known" Spam sites (a black list) and then block any incoming email from those addresses. A white list is the opposite of a blacklist. It identifies those sites from which an email server may receive SMTP requests.

IP blocking can occur anywhere along the network path beginning with the point of the origin and ending with the destination server. The message can be stopped before it is sent by blocking a DNS query from known Spam sites. If the Spam site is unable to

query the DNS MX record for a particular domain then it is unable to send the message. The more common practice is to have the email server (or a email routing server situated in front of the email server) validate the IP address of the message source from a "known" list of Spam (or Open Relay) sites. If the message source is on the list of sites to be blocked, then the SMTP connection is terminated.

There are many problems with IP blocking as a form of Spam control. The problem with a whitelist is that in order to receive email from a new site, the systems administrator must manually add the site to the "allowed" email list. There are several problems with blacklists, not the least of which is that it is an extremely labor intensive task to adequately research the source of the Spam. Often Spam is routed through several open SMTP relay servers to disguise its origin. For blacklists to be effective they must also include open relay sites. This means that if a systems administrator makes an email server configuration error and doesn't correct it prior to being placed on a blacklist, then all email coming from their site, including legitimate email, is subject to being blocked by those that subscribe to the blacklist. To add to their already daunting task of researching new IP addresses to block, the researchers who maintain the blacklist must also do follow-up auditing to remove newly compliant sites from their blacklist. It is very easy to be blacklisted but very difficult (and costly) to be removed from the list.

The Mail Abuse Prevention System (MAPS) is one of the primary proponents for the use of the blacklist as a method for Spam control. They call their list the Realtime Blackhole List (RBL). "But the accuracy of the MAPS' blacklist has been called into question... David Nelson, a senior industry analyst at Giga Information Group, says a recent study found that Brightmail, a for-profit blacklisting and filtering service, blocks 94% of spam with 1% false positives. However, MAPS was found to block 24% of spam with 34% false positives." [17] In addition to concerns about an inordinate number of false positives with MAPS RBL there are also concerns that their blacklist inclusion criteria is too broad and thereby targeting too many legitimate internet users and businesses. There are other providers of blacklists. ORBZ and SPEWS also provide their own lists, and each of these organizations has also been criticized for inequities in the management of their respective blacklist.

How effective are Spam filters?

Early Spam lexical filters used a set of rules to recognize specific Spam features. A good example of this is the keyword filter. Simply adding the word "click" to a Spam filter apparently can catch close to 80% of Spam messages with only 1.2% false positives. [18] The latest technology in spam filters is the "statistical" filter". "Statistical filters look at the entire content of each incoming email and decide whether or not it's Spam based on its overall similarity to previous Spams. This new kind of filter routinely catches over 99% of current Spam with near zero false positives." [19] To be effective the statistical filters require training.

Statistical filters work by analyzing a user's Spam messages and also their non-Spam email messages. The filter creates a word frequency index from each message category. It then creates an index value based on the ratio of Spam to non-Spam for any given word. This type of statistical filter is called Bayesian filtering.[a] Paul Graham in his web article, "A Plan for Spam" makes a convincing argument for using the Bayesian filter.

> But the real advantage of the Bayesian approach, of course, is that you know what you're measuring. Feature-recognizing filters like SpamAssassin assign a spam "score" to email. The Bayesian approach assigns an actual probability. The problem with a "score" is that no one knows what it means. The user doesn't know what it means, but worse still, neither does the developer of the filter. How many *points* should an email get for having the word "sex" in it? A probability can of course be mistaken, but there is little ambiguity about what it means, or how evidence should be combined to calculate it. Based on my corpus, "sex" indicates a .97 probability of the containing email being a spam, whereas "sexy" indicates .99 probability. And Bayes' Rule, equally unambiguous, says that an email containing both words would, in the (unlikely) absence of any other evidence, have a 99.97% chance of being a spam.
>
> Because it is measuring probabilities, the Bayesian approach considers all the evidence in the email, both good and bad. Words that occur disproportionately *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. [20]

Paul Graham's evangelistic tone aside, the statistical approach to spam filtering using Bayesian methodology appears to have an edge on other, less sophisticated methods. Rules, keyword and IP blocking based systems are not flexible enough to handle the ever-changing signature of the common Spam message. Further, approaches other than the Bayesian approach do not take into account the personal preferences of the email user, whereas the Bayesian analysis factors in each user's legitimate message content. For this reason the Bayesian filter is the most user-oriented anti-Spam approach available today.

There are several issues with statistical filters. One of the major issues is that they need to be trained and this involves effort by the end user. The user must not only identify messages that constitute Spam, but also fix incorrect message classifications. For these reasons server-level filters may be difficult to implement.

Other efforts:

[a] Ten days ago I installed a Bayesian Filter called POPFile (http://popfile.sourceforge.net/). During this period the filter processed 892 email messages and had 26 classification errors resulting in an accuracy of 97% and a 3% false-positive rate and dropping daily. Of the 892 emails 87% are Spam.

The Anti-Spam Research Group (ASRG) is an organization of professionals that are taking a different approach to Spam control. Unlike MAPS approach or end user filtering, ASRG is investigating the problem as a large-scale network issue. They are in the early stages of creating a design specification that addresses the Spam problem through the implementation of a consent-based architecture. [21] There are three primary components to this architecture:

1) Consent Expression Component: where the recipient positively expresses a consent or non-consent policy for receiving certain types of communications.

2) Policy Enforcement Component: where the systems (or sub-systems) within the internet network architecture are empowered to enforce the recipients stated consent policy. This can follow two main logical paths: "fail-open" or "fail-closed". In the "fail-open" gate the messages that do not have consent are filtered. In the "failed-closed" gate only the messages that have consent are passed to the recipient.

3) Source Tracking Component: where the message source is more readily identifiable within the network.

It should be noted that ASRG's work is in the early stages and no proposals have yet to surface.

What to do to?

1) First and foremost is not to contribute to the problem. Make sure that each server on the network is protected from unauthorized SMTP access. If SMTP services are not required on the server, then the best practice is to shut down this service. On servers that do require SMTP services, then restrict access to SMTP to only locally hosted domains.

2) Implement and maintain a network-centric, managed anti-virus program for all workstations and email servers.

3) Decide if IP blocking will work in your organization. The false-positive rate may be acceptable for some companies with a fairly limited corporate email universe. If IP Blocking works and the email server's MTA supports it, then configure the server to utilize one of the available black lists.[22] Don't forget to identify known "good" sites within the server's "white-list" configuration.

4) If IP blocking will not work in your organization, then subscribe to an email filtering service such as Brightmail.

5) If options 3 or 4 are unacceptable, then implement email client filtering at the workstation level and educate users on how to "train" the filters. This also may

be used as part of a layered defense against Spam. (see Paul Graham's list of Spam filters for further information. [23])

Conclusion

Just as a layered defense system is important for network security, a multi-layered defense against Spam is also useful. Fundamentally, Spam is a commercial enterprise and as such, if it loses its effectiveness in the marketplace there will be less profit incentives for the use of Spam. Less profit means less commercial Spam. This does not account for the spammers who use it as a delivery method for a virus/Trojan payload. Programmers who write malicious software will always find a way to deliver their software to the most vulnerable systems.

Each layer of Spam defense has inherent strengths and weaknesses. Because Spam is difficult to define and because its point of origin is often outside of US jurisdiction, it is difficult to write legislation to mitigate the problem. Legislation, however ineffective, is an important first step towards defining the problem. Proper legislation can also institute fairness into the methods that are used to fight Spam. On the technical front, until anti-Spam software can achieve a zero false positive rate, users will need to monitor their own messages for email that was incorrectly identified as Spam. Statistical filters can be applied either at the email destination server or integrated into the email client at the user's workstation – or better yet, both. ASRG's policy driven, network-centric Spam control approach seems promising, but a solution such as this is further down the road and not imminent. Completely eliminating Spam is an unrealistic goal, but reducing its effectiveness as a marketing method and protecting end users from its overwhelming presence is more realistic.

[1] "Mail Abuse Prevention System, Definition of 'spam' ".
URL: http://www.mail-abuse.org/standard.html (February 25, 2003).

[2] Graham, Paul. "A Plan for Spam". August 2002.
URL: http://www.paulgraham.com/spam.html (March 2, 2002).

[3] Graham, Paul. "A Plan for Spam". August 2002.
URL: http://www.paulgraham.com/spam.html (March 2, 2002).

[4] "KnowledgeShare - White Papers, Spam". Vicomsoft.
URL: http://www.vicomsoft.com/knowledge/reference/spam.html (February 25, 2003).

[5] Weinschenk, Carl. "Prevention is the Best Medicine: Creating an Anti-Spam Strategy". ServerWatch. January 9, 2003. URL:
http://www.serverwatch.com/tutorials/article.php/1567361 (March 2, 2003).

[6] Weinschenk, Carl. "Prevention is the Best Medicine: Creating an Anti-Spam Strategy". ServerWatch. January 9, 2003. URL:
http://www.serverwatch.com/tutorials/article.php/1567361 (March 2, 2003).

[7] "The Economics of Spam". ePrivacy Group. (no date provided).
URL: http://www.eprivacygroup.com/article/articlestatic/58/1/6 (March 14, 2003).

[8] "Spammers Send Over One Million Unwanted Valentine's Day Offers This Year". Brightmail. February 14, 2003. URL:
http://www.brightmail.com/pressreleases/021403_Valentine.html (March 3, 2003).

[9] Rutherford, Paul. "Bugwatch: Beware the Valentine's Day Massacre". February 13, 2003. URL: http://www.vnunet.com/News/1138770 (February 26, 2003).

[10] Tuesday, Vince. "Spam Issue Viewed As IT Failure". ComputerWorld. January 13, 2003. URL:
http://www.computerworld.com/securitytopics/security/story/0,10801,77396,00.html (February 25, 2003).

[11] Gaudin, Sharon. "Spam Explodes Over Holidays -- Up 1,000%". December 31, 2002.
URL: http://www.internetnews.com/stats/article.php/1562611 (March 3, 2003)

[12] "Spam Laws: United States: Federal Laws".
URL: http://www.spamlaws.com/federal/index.html (March 15, 2003).

[13] "Spam Laws: United States: State Laws: Summary".
URL: http://www.spamlaws.com/state/summary.html (March 15, 2003)

[14] "Spam Laws: European Union/EEA". URL: http://www.spamlaws.com/eu.html (March 15, 2003)

[15] "Spam Laws: Other Countries". URL: http://www.spamlaws.com/world.html (March 15, 2003)

[16] Graham, Paul. "Filters vs. Blacklists". September 2002. URL: http://www.paulgraham.com/falsepositives.html (March 2, 2002).

[17] Gaudin, Sharon and Gaspar, Suzanne. "The Spam police". Network World. September 10, 2001. URL: http://www.nwfusion.com/research/2001/0910feat.html (March 2, 2003).

[18] Graham, Paul. "A Plan for Spam". August 2002. URL: http://www.paulgraham.com/spam.html (March 2, 2002).

[19] Graham, Paul. "Will Filters Kill Spam?". December 2002. URL: http://www.paulgraham.com/wfks.html (March 3, 2003).

[20] Graham, Paul. "A Plan for Spam". August 2002. URL: http://www.paulgraham.com/spam.html (March 2, 2002).

[21] "Anti-Spam Research Group (ASRG)". URL: http://www.irtf.org/charters/asrg.html (March 17, 2003).

[22] "Controlling e-mail spam, Blocking by MTAs". URL: http://spam.abuse.net/adminhelp/mail.shtml (March 14, 2003).

[23] Graham, Paul. "Spam Filters". URL: http://www.paulgraham.com/filters.html (March 14, 2003).