Spam in Email Inboxes

Ursula Dantin **John Paynter**

Department of Information Systems and Operations Management
The University of Auckland
Auckland, New Zealand
u.dantin@auckland.ac.nz

ABSTRACT

Since 2002, an exponential increase in spam and spam-related virus attacks could be observed. The practice of spamming will only decrease when the financial incentive to spam is minimised. Spam can only be reduced by a consolidated worldwide effort, combining legal measures of individual countries and international collaboration of enforcement, as well as spam blocking technology, and education on best practices in companies regarding mass emailing. The complexity and cost of this task means the spam phenomenon will be an ongoing issue for some years to come. As most reports on spam originate from the media or service providers we present the results of a small survey to see if it is the problem that is reported elsewhere.

Keywords

Spam, email, filters, legislation

1. INTRODUCTION

During 2003, an escalating discussion in the media could be observed on the issue of unwanted commercial mass emailing over the Internet, generally referred to as spam. According to media reports, more and more people seemed to complain about spam flooding their email accounts. The reports suggested a serious spam problem on a worldwide scale. What evidence existed about the real extent of the phenomenon, and what were the most successful strategies suggested in the IS media for combating it?

After summarising the facts on spam as portrayed in the media before mid of 2004 and from the small sample of an email user survey conducted in late 2003, this article looks at anti-spam software, how it works, and summarises what strategies were recommended to fight spam. Legislative moves in major industrialised countries by early 2004 are noted, and how these might impact spam proliferation. After an outlook at spam related issues past mid 2004, conclusions are drawn about the state of the phenomenon spam in early 2004.

2. PRELIMINARY SPAM SURVEY

An online survey was conducted late in 2003 from The University of Auckland, asking email users to comment on their level of spam received and whether they experience it as a real problem. Technical problems with the survey resulted in only about 30 responses (mainly from the IT sector in New Zealand) being collected, most of them incomplete. However, the results can give a first indication on email user sentiments on spam to compare with the reporting in the media.

3. FACTS ON SPAM

3.1 How to Get Spammed

Although some survey participants showed knowledge of address harvesting and selling practices, they were surprised by the types of products advertised to them.

Already in 2002, a group of researchers at the Centre for Democracy and Technology in the U.S. found that of the 250 artificial email addresses they created and posted in public and corporate domains, only addresses in spambot-readable form received junk mail. Of the 8,400 spam messages received (from a total of about 10,000 messages over 6 months) 97 percent were delivered to addresses that had been posted on public domains linked to major portals such as Yahoo and AOL (Ingham, 2003). This was a

clear indication that at least at that time, spammers used address-harvesting software to scan public domains.

Research in Great Britain suggested that the increase in always-on broadband connections and the resulting increase in Internet use was at least partially responsible for the proliferation of spam (McCue, 2003). The high data capacity of broadband made downloading so much faster that people surfed more, leaving their email addresses in many places. PCs with always-on broadband were of interest to spammers and hackers as a worthwhile target for hijacking as an unwitting host (Trojan horse virus) for future spam or virus proliferation.

From 28 survey participants two did not receive any spam (one because of massive use of anti-spam protection, and one because he admitted to being a spammer himself); one third received between 1-5 and 5-50 spam emails daily. Three received more than 50 per day. These were roughly equally accessing their email at work or through their own private ISP connection. Still, it seemed that overall most spam was sent to corporate addresses – over 70 percent in July 2003 according to Postini, a spam filter service provider (Gifford, 2003a).

Direct marketing associations claimed there was a place for legitimate mass mailing. They used increasingly emailing for its extremely low cost compared to conventional mail. Even mailing to purchased email address lists was considered legitimate, if the receiver was offered to opt-out of future emailing.

3.2 Content of Spam Messages

The statistics on the content of email spam messages from two spam filter providers showed that in September 2003, most were product offers (20 to 30 percent), followed closely by so called 'adult content' and healthcare (around 20 percent each). Scams were under 10 percent. (Gifford, 2003b, Clearswift, 2003).

A question regarding spam content was not in the survey because it would often be impossible to tell the exact content without viewing. Since people tried hard to ignore spam, their estimate on spam content by percentage would most likely be inaccurate.

Viruses spread by spammers had mainly one of two aims: harvesting more email addresses from address books or corporate customer databases, or planting a Trojan horse type virus that turned the computer into a proxy server sending more spam when connected to the Internet. This could result in normal email users being included in blacklists used by spam filter providers to identify spammers.

3.3 Why Spam

It did not seem unusual for spammers to send up to 200 million messages a day. At this rate, even one hit in half a million was considered significant (Gifford, 2003b) and made the effort financially viable. Asked in the survey how often they would reply to spam, from 28 answers most said 'never'. Only one each said 'often' and 'sometimes', three said 'rarely'.

Direct marketers using e-marketing were most interested in the conversion rate, meaning the percentage of actual sales achieved. DoubleClick, an Internet direct marketing service, still reported a conversion rate of 2.68 in 1,000 emails sent (3rd Quarter 2003), resulting in an average return of US$0.26 per email sent (Rodgers, 2004). This compared well to the average cost per email of a fraction of a cent.

An explosion of virus attacks via spam messages could be observed during late 2003 and early 2004 in statistics published by some ISPs. It was probably related to the criminal spammers, who were fighting the use of better spam and virus filters by setting up Trojan horse hosts. Apart from utilising the hosts processing capacity to exponentially increase the number of emails sent, this was fooling spam filter blacklists by disguising the point of origin of the message. Despite the hurdles the anti-spam software created, the cost to spammers was still not high enough to make the business of spamming unprofitable.

3.4 Cost of Spam

A spammer could pay as little as US 0.025 cent to send an email message (Hansell, 2003). The real cost was with the recipient. The most significant, and at the same time hardest to specify, was the cost of lost productivity: deleting spam and finding lost genuine messages; plus the cost of lost business from false positives, genuine emails that were destroyed as spam. On top of this came the additional time employees spent dealing with

their own posted legitimate emails that were rejected by the recipient's spam filters. Other more obvious costs were processing and storage costs of spam messages for the recipient and the ISP, as well as the cost of trying to detect and stop spam, plus building and constantly upgrading the filter software. Even though these costs were almost as tiny per single message as the sending cost to the spammer, the sheer volume of such messages sent each day added up to substantial amounts in terms of economic impact. The highest cost estimate in 2003 came from Nucleus Research, a commercial advisory service, forecasting the total economic cost of spam at US$87 billion in the U.S. for 2003 (Hansell, 2003).

The cost structure was similar for ISPs and network administrators. The actual volume in Internet traffic caused by spam was just a fraction of total Internet traffic. Because spammers paid for network capacity used, they had a financial incentive to keep messages small. However at MCI, one of the largest American backbone carriers, unpaid bills from evicted spammers were the single highest cost in mid-2003. At ISPs, like at the receiving end, a substantial cost was hidden in the human element. MCI was reported to have received half a million complaints a month regarding spam transmitted through their network. This meant MCI had employees solely investigating complaints and disconnecting offenders (Hansell, 2003).

In the survey we asked 'Does spam bother you?'. From 28 responses, 12 were 'somewhat significant' and 6 'significant'; together two thirds of the respondents. They seemed slightly more concerned about loosing genuine messages, their overflowing inboxes, and the waste of time than about disturbing content.

As a last item, the cost of scam fraud should be considered. Phishing for example was on the rise since 2003. Here spammers pretended to be a big company or major bank trying to extract personal details like account numbers, passwords, and credit card details from unsuspecting individuals. Criminals used these details to empty bank accounts and so directly harm the fraud victim. Luckily, banks limited the liability of their customers in most cases. The more substantial cost was to businesses losing money from purchases if the victim could prove the fraud.

4. HOW TO FIGHT SPAM

According to the media, the exponential increase in spam emails experienced during 2002 and 2003 resulted in end-user pressure to act on systems administrators, and ultimately ISPs. Figures published by various anti-spam providers claimed that in September 2001, only 8 percent of all commercial email was unsolicited. But by July 2003 it was over 50 percent (Brightmail quoted in Gifford, 2003b). An increasing number of commercial anti-spam service providers could be seen as a clear indication for a real (or only perceived?) demand in this area of Internet services. In the survey, 19 respondents of 28 said they did not use any spam protection, most (8 cases) because they 'can't be bothered'. However, it was questionable whether they were not already protected anyway to some degree by the automatic spam scanning at their ISP.

4.1 How Do Spam Filters Work

One of the more sophisticated methods of spam filtering used Bayesian logic, calculating from the number of times an event had not occurred the probability that it will occur in future. Applied to spam filters this was used to "train" recognition of spam into the filter software. Commercial spam filter software using this method claimed a false positive rate of only one message in 90,000 (Gifford, 2003a).

Other tools scanned the content of incoming emails for heuristic triggers or looked at what the email asked the recipient to do. Since 90 percent of spam carried an embedded URL the scanning software looked for those, confident that most were actually spam (Brislen, 2003). Alternatively, the software looked at the absence of something that a normal email would contain. Some popular open source filters used the absence of a signature as an indicator of spam. Unfortunately, spammers then fooled the software by simply adding a signature to their message (Lemos, 2003).

More conventional methods that were sufficient for most users before spam occurred in such high daily volumes are still in use. The most widely used were blacklists and whitelists. Blacklists of sender addresses that should be blocked required constant updating to take into account the changing addresses spammers use. In early 2004, most spam used forged return

addresses (Hansell, 2004). Since spammers increasingly forged the addresses of innocent users and used Trojan horse viruses to send their spam, blacklists were becoming increasingly unreliable and could actually do more harm than good.

Whitelists contained addresses of senders from whom the email recipient explicitly wished to receive mail from, even if they were mass mailings. Greylisting, a method used to create a customised whitelist for a recipient, initially returned every email from a non-listed sender asking for a re-send as validation. This needed to happen only once, after which the sender was added to the whitelist for future mailings. Spammers using robots to send their mailings could not recognise and process these requests for reply, especially when the sender address used was a fake. A response mechanism could increase the cost for spammers to an uneconomic level (Harris, 2003).

Brightmail, who filtered in 2003 by its own account about 11 percent of email traffic worldwide (Gifford, 2003b), combined more conventional scanning methods with an additional reliance on the human brain as the most accurate and flexible spam detection tool. The company asked their clients, mostly ISPs and larger corporates, to host decoy email addresses on their networks that are never actively used. All incoming emails from these about 1 million decoy addresses worldwide were by definition unsolicited. Humans made a final decision and the spam filters were updated every few minutes to stay highly effective. Brightmail claimed that, thanks to this unique addition to conventional scanning methods, they achieved a very low false positives rate of one in a million (Brislen, 2003).

4.2 Technical Anti-Spam Measures

In New Zealand two ISPs were reported to have implemented commercial anti-spam filters in addition to their already existing anti-virus filters; TelstraClear in September and Xtra in November 2003. In the U.S., Yahoo! launched new anti-spam tools for its web-based email services at the end of October 2003.

TelstraClear reported that the filters blocked between 30 and 60 percent of messages in the first few days (Gifford, 2003b). This indicated that TelstraClear had invested in twice the server infrastructure and bandwidth necessary, incurring double the licensing cost, not to mention the cost of dealing daily with hundreds of customers upset about spam they received. The intercepted emails were stored initially for one month, enabling end users to retrieve any false positives during this time. AOL reduced the cost of storage even further by storing identical spam emails only once and giving all recipients of that email access to this one copy (Hansell, 2003). However, the exponential increase in spam and spam related viruses observed in late 2003 and early 2004 was very likely a direct result of the installation of these filters by all majors ISPs and could quickly eliminate the initial reduction of spam in email inboxes. Each time spam filters were improved, spammers found a work-around.

The best results in spam filtering could be achieved by a combination of methods: Blacklists stop known spammers, greylists eliminate spam sent by spambots, decoy email boxes alert to new spam attacks, and Bayesian filters detect spam and virus attack emails right from the start, even before the blacklist is updated. False positives were still a problem, along with a new kind of spam, bouncing emails created by viruses (Bell, 2003).

4.3 Human Anti-Spam Measures

Since statistics showed that in early 2004 most spam was still directed at organisations and not at private email addresses, an important step towards minimising unsolicited and unwanted emails was raising the awareness of employees. Many people were still not fully aware of the information flood they were able to generate via their email facility for customers and colleagues alike.

Usually, spam watchdog organisations and ISPs took on the task of publishing blacklists. Spammers tried of course to make their identification as difficult as possible. If a spammer's mail server finally was shut down, ISPs or compilers of blacklists sometimes even got phone calls threatening to kill from hardcore criminal spammers (Sturgeon, 2003a).

Getting an ISP to actually shut down a spam server was notoriously difficult in countries such as Russia, China, and other parts of Asia, as well as in some South American countries. The lack of relevant laws or law enforcement in these regions had attracted a large number of

spammers. However, the real perpetrators usually lived in the U.S. and wanted their spam to be sent to U.S. consumers. In September 2003, the Spamhaus Project, a British anti-spam group collating blacklists, estimated that 100 of North America's most prolific spammers were situated in the suburbs of Beijing (Sturgeon, 2003b).

A more unconventional way to combat spam, a kind of guerrilla warfare, was used by some not-for-profit anti-spam lobby groups. They tried to harass notorious spammers by publishing their contact details, usually on the web, for everybody to send unsolicited mail. An example where this strategy seemed to work was the case of a local penis enlargement spammer, exposed in New Zealand in August 2003 by the country's largest daily newspaper, the NZ Herald (Saarinen, 2003). After a few days the spammer publicly announced that he had stopped doing business, pleading for the harassment of his family to stop.

4.4 Legal Anti-Spam Measures

Efforts to control the spam flood with legal measures seemed to be not very successful by early 2004. Even though the U.S. enacted a federal anti-spam law from 1 Jan 2004, the IS community expected no significant improvement from it. This CAN-SPAM bill was overruling more restrictive and effective state laws that were for example previously in place in California. Some commentators even suggested that the bill legalised unacceptable spamming practices (McNamara, 2004) because lawmakers bowed to the direct marketing lobby.

In Europe, national and EU wide efforts in early 2004 seemed to favour the opt-in approach where recipients must expressly declare their willingness to receive the direct marketing material for spam to be legal. Even though the anti-spam lobby was more favourable of this approach, practice showed that people were often tricked into giving consent without realising it. In the South Pacific, Australia was ahead of New Zealand introducing a bill in 2003 while New Zealand lawmakers were still assessing their options in early 2004.

In reality, all national laws hardly affected hardcore spammers. They were an international phenomenon, often closely associated with criminal entities such as the Russian mafia. National laws could only be applied if the offender lived within the jurisdiction of that law. It was clear that in the U.S. most of the companies buying the services of a non-U.S. spammer were actually situated in the U.S., also shipping their wares from the U.S. But it was very difficult and time consuming, and therefore expensive, to collect enough evidence about the out-of-state spamming activities to prove the connection and prosecute successfully. One major criticism of the CAN-SPAM bill was the total lack of resources initially allocated to this task (Caruso, 2004).

5. OUTLOOK

The next major development in terms of spam was expected to be spam to mobile Internet devices such as cell phones and PDAs. One concern was the threat this posed to children possessing these devices being exposed to explicit pornographic material (Saarinen, 2003).

Major ISPs experienced ongoing problems with overloading of servers caused by the exponentially increasing amount of spam. On a hardware level, this was starting to be counteracted by the installation of different servers that could process high volumes of small files and detect spam already at this point.

As a result of the preliminary survey, the literature review and current trends a revised survey will be published on the web. The results of which will be presented at a later date.

6. CONCLUSION

All figures published in the media on the amount of spam circulating and the financial and psychological impact of this phenomenon seemed to come from sources with a commercial interest in spam fighting. As they might have a vested interest in overstating the problem, these figures should be treated with caution. First responses to an end-user survey into the spam phenomenon seemed to indicate that people were less annoyed by spam than media reports suggested. A larger survey will hopefully help to give a more rounded view.

The explosion of spam and spam-related virus attacks in late 2003 and early 2004 seemed clearly related to the introduction of commercial anti-spam filters by all major ISPs. In an effort to retain their profits, spammers were fighting back by increasing the amount of spam emails

sent. Still unprotected users could be forced to introduce spam protection tools.

The development of any legal tools in the fight against spam seemed very much bogged down in red tape and political considerations, not to mention the problems resulting from the fact that big volume spammers usually operated across borders and could keep outside the reach of any threatening jurisdiction.

Many experts in the IT area suggested that the spam problem would not disappear any time soon. A combination of different technologies and strategies will be in the long term most effective in defeating spam. "There's no silver bullet. You need laws to deter spammers. You need technology to block spam and you need the direct marketers to adopt best practices so they avoid sending spam in the first place." - Enrique Salem of Brightmail (Brislen, 2003). A major challenge seemed to be narrowing the time gap between the start of a spam or virus attack and the moment the filters were updated to recognise these emails and stop them. As far as commercial spam filters were concerned, false positives were still a problem even for major ISPs, like AOL, that could afford the best products.

Media reports considered that removing the financial incentive was the only way to effectively break the tide of spam, by preventing most spam emails from reaching a consumer. Gullibility seems to be part of human nature. This means that as long as a certain percentage of emails get through to the consumer, inevitably some people will reply and the spammer has the chance to make his cut. This practice will only cease once the profits disappear.

REFERENCES

Bell, S. (2003) "Filters causing rash of false positives: TelstraClear's new virus and spam screening service gets mixed reviews". Posted Sep 26. Accessed Oct 30, 2003. <http://computerworld.co.nz/webhome.nsf/printdoc/FD6838027D502EB6CC256DAC000D90D4!opendocument>

Brislen, P. (2003) "Spam epidemic has three years left to run, says anti-spam head. 'There's no silver bullet'". Posted Sep 06. Accessed Oct 30, 2003. <http://computerworld.co.nz/webhome.nsf/UNID/A4F963DCA80A4B80CC256D96001A6ED4!opendocument>

Caruso, J. (2004) "Anti-spam law is a start, panel says". Network World, Jan 05. Accessed Feb 12, 2004. <http://www.nwfusion.com/news/2004/0105spam.html>

Clearswift (2003) "Increase In spyware proves spammers always on the look out". Posted Oct 06. Accessed Oct 30, 2003. <http://www.clearswift.com/news/pressreleases/items/216.asp>

Gifford, A. (2003b) "Way to fight spam is to take away commercial gain". NZHerald, Oct 03. Accessed Oct 30, 2003. <http://www.nzherald.co.nz/storydisplay.cfm?storyID=3526731&thesection=technology&thesubsection=general>

Gifford, A. (2003a) "Old minister helps beat spam". NZHerald, Jul 29. Accessed Aug 04, 2003. <http://www.nzherald.co.nz/storydisplay.cfm?storyID=3515081&thesection=business&thesubsection=technology&thesecondsubsection=information>

Hansell, S. (2004) "Gates backs e-mail stamp in war on spam". New York Times, Feb 02. Accessed Feb 03, 2004. <http://www.nytimes.com/2004/02/02/technology/02spam.html>

Hansell S. (2003) "Diverging estimates of the costs of spam". New York Times, Jul 28. Accessed Jul 29, 2003. <http://www.nytimes.com/2003/07/28/technology/28SPAM.html>

Harris, E. (2003) "The next step in the spam control war: Greylisting". Accessed Apr 10, 2005. <http://projects.puremagic.com/greylisting/whitepaper.html>

Ingham, R. (2003) "I think, therefore I'm spam". Posted Jun 25. Accessed Oct 30, 2003. <http://cooltech.iafrica.com/features/248144.htm>

Lemos, R. (2003) "Signed spam tries to fool filters". Posted Oct 13. Accessed Oct 30, 2003. <http://www.silicon.com/research/specialreports/thespamreport/0,39025001,10006378,00.htm>

McCue, A. (2003) "UK fed up with its fat pipes". Posted Oct 28. Accessed Oct 30, 2003. <http://www.silicon.com/networks/broadband/0,39024661,39116646,00.htm>

McNamara, P. (2004) "Primary purpose? . . . To deceive". Network World, Feb 09. Accessed Feb 12, 2004. <http://www.nwfusion.com/columnists/2004/0209buzz.html>

Rodgers, Z. (2003) "Cleanliness key to e-marketing success". Posted Dec 05. Accessed Jan 04, 2004. <http://cyberatlas.internet.com/markets/advertising/article/0,,5941_3285591,00.html>

Saarinen, J. (2003) "Spammer ducks for cover as details published on web". NZHerald, Jul 19. Accessed Aug 19, 2003. <http://www.nzherald.co.nz/storydisplay.cfm?thesection=news&thesubsection=&storyID=3518682>

Sturgeon, W. (2003b) "Major victory in war on spam". Posted Sep 09. Accessed Oct 30, 2003. <http://www.silicon.com/research/specialreports/thespamreport/0,39025001,10005930,00.htm>

Sturgeon, W. (2003a) "Spam Summit: Laws flawed from outset - but still a step in the right direction?". Posted Jul 01. Accessed Oct 30, 2003. <http://www.silicon.com/research/specialreports/thespamreport/0,39025001,10004937,00.htm>