Trends in Spam Products and Methods

Geoff Hulten – Anthony Penta – Gopalakrishnan Seshadrinathan – Manav Mishra
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
{ghulten, apenta, gopalaks, manavm}@microsoft.com

Introduction

In this paper we analyze a very large junk e-mail corpus which was generated by a hundred thousand volunteer users of the Hotmail e-mail service. We describe how the corpus is being collected and then discuss how both the products being advertised by spam and the specific exploits being used to avoid spam filters have changed over time.

Every day we randomly select one message from the mail stream of each Hotmail volunteer and ask that user to classify it for us. Thanks to these users, we have been receiving tens of thousands of hand classified messages per day, every day for the past year – our database currently contains over ten million classified messages. In this paper we further analyze two samples of the spam from this data, one from early 2003, and one from early 2004. We categorized the spam by the type of product it is selling, and by the types of exploits it uses to avoid spam filters.

We are aware of very few other large scale studies of spam. One is the FTC report on false claims in spam [1]. Our study differs by using data sets that were created by randomly sampling over the entire mail stream, rather than by relying on users to report e-mail that offended them; by reporting changes in spam data over time; and by reporting on more categories of spammer exploits. Another relevant large scale study is our analysis of the geographic origins of spam [2].

Trends in Spam

We randomly sampled 1,000 spam messages that arrived at Hotmail between 2/1/04 and 3/1/04, called the *2004 Spam* set in what follows. We also sampled 200 messages that arrived at Hotmail between 3/15/03 and 4/15/03, called the *2003 Spam* set in what follows. We hand-examined each of these messages to determine what type of product they were promoting and what exploits they use. Table 1 shows the product information. By far the largest increase over the last year has been in Porn/Sex Non-graphic, and the majority of this increase has been in products selling sexual enhancers.

Table 1 : Trends in Categories of Spam

Product	2003 Spam	2004 Spam	Delta (Absolute%)	Description
Porn/Sex Non-graphic	17%	34%	17%	Enhancers with sexual connotation, links to porn.
Insurance	1%	4%	3%	Health, dental, life, home, auto insurance.
Rx / Herbal	8%	10%	2%	Cheap drugs or herbal supplements.
Financial	12%	13%	1%	Refinancing, get out of debt, financial advice.
Travel / Casino	2%	3%	1%	Selling airlines tickets, hotel reservations, rental car. Internet casino sites. Other gaming sites.
Scams	8%	6%	-1%	Get rich quick, Phisher scams, etc.
Newsletters	9%	6%	-3%	Any newsletter that isn't selling something.
Other Spam	13%	8%	-5%	Everything else that appears to be spam.
Porn/Sex Graphic	13%	7%	-5%	Anything that contains pornographic images.
Dubious Products	20%	10%	-10%	Pirated software, diplomas, etc.

Table 2 shows the exploits that are present in the messages in the samples. Text Chaff and Domain Spoofing are the two most popular exploits in our data, and these have been used fairly consistently over time. There have been large increases in many categories of exploits. Most notably Word Obscuring and URL Spamming were not used much in 2003 but are starting to see widespread use in 2004. Only two exploits declined, namely URL Obscuring and Character Encoding. We believe this is because these exploits are relatively easy to detect using simple regular expressions, spam filtering products are beginning to reliably catch messages that use them, and so spammers are reacting. Another interesting fact is that the nature of text chaff has changed drastically: in 2003 it was almost all strings of random characters, in 2004 about a quarter of it is made up of actual words.

Table 2 : Trends in Exploits in Spam

Exploit	2003 Spam	2004 Spam	Delta (Absolute %)	Description
Word Obscuring	4%	20%	16%	Misspelling words, putting words into images, etc.
URL Spamming	0%	10%	10%	Adding URLs to non-spam sites (e.g. msn.com).
Domain Spoofing	41%	50%	9%	Using an invalid or fake domain in the from line.
Token Breaking	7%	15%	8%	Breaking words with punctuation, space, etc.
MIME Attacks	5%	11%	6%	Putting non-spam content in one body part and spam content in another.
Text Chaff	52%	56%	4%	Random strings of characters, random series or words, or unrelated sentences.
URL Obscuring	22%	17%	-5%	Encoding a URL in hexidecimal, hiding the true URL with an @ sign, etc.
Character Encoding	5%	0%	-5%	Phar&#109;acy renders into Pharmacy.

The average spam message from 2003 had 1.33 exploits and the average spam message from 2004 had 1.73 exploits – an increase of 0.40 per message. Table 3 shows how exploit usage has varied by product. One interesting point is that there has been a large decrease in the number of exploits being used in graphic pornographic messages. Scams, Financial, and non-graphic porn have seen the largest increases.

Table 3 : Trends in Number of Exploits per Message

Category	2003 Spam	2004 Spam	Absolute Delta
Scams	1.07	1.89	0.82
Financial	1.26	1.88	0.62
Porn/Sex Non-graphic	1.85	2.44	0.58
Travel / Casino	0.75	0.98	0.20
Other Spam	0.76	0.79	0.32
Insurance	1.50	1.52	0.02
Newsletters	0.00	0.00	0.00
Rx / Herbal	2.13	2.12	-0.01
Dubious Products	1.20	1.15	-0.05
Porn/Sex Graphic	2.16	1.33	-0.83

Summary

In the past year spam messages selling Non-graphic porn/sex have increased dramatically. Spammers have reduced their reliance on URL Obscuring and Character Encoding, while increasing their use of URL Spamming and Word Obscuring. Spammers also seem to be working harder: today the average spam message contains more exploits than it did a year ago.

Acknowledgements: Thanks to Robert Rounthwaite and Joshua Goodman, and thanks to the Hotmail team for setting up the Feedback Loop, especially Pablo Stern, Eliot Gillum, and Raj Pai.

References

[1] FTC Division of Marketing Practice. False claims in spam. http://www.ftc.gov/opa/2003/04/spamrpt.htm. April 30, 2003.

[2] Hulten, G., Goodman, J., and Rounthwaite, R. Filtering spam e-mail on a global scale. In Proceedings of the Thirteenth International World Wide Web Conference (pp. 366 - 367), 2004. New York, NY: ACM Press.