A New Approach to Defeating Spam

An Osterman Research White Paper

Published March 2008

Osterman Research, Inc. • P.O. Box 1058 • Black Diamond, Washington 98010-1058
Phone: +1 253 630 5839 • Fax: +1 866 842 3274 • info@ostermanresearch.com • www.ostermanresearch.com

Overview

Junk postal mail is a nuisance for those who receive it, but it is limited by two important economic factors: a) junk mail costs something to produce and, as a result, b) senders of junk mail must achieve acceptable content-to-customer conversion rates in order to make the sending of their information economically worthwhile.

The electronic equivalent of junk postal mail – spam – however, operates under no such economic constraints. Hundreds of millions of spam messages can be sent for a minimum investment and conversion rates can be extraordinarily low for spammers to turn a sizable profit. In fact, spammers can also make money without having to sell anything simply by directing people to Web sites that contain advertising messages. Further complicating the problem for recipients of spam is the fact that spammers are continually developing newer and more innovative techniques to defeat conventional spam-filtering technologies. This results in greater storage requirements, constrained bandwidth, reduced employee productivity and a host of other problems.

> What is needed…is a system that can defeat spam reliably, generate a minimum of false positives, will not be rendered obsolete by new spammer techniques and be deployed and managed inexpensively.

What is needed, therefore, is a system that can a) defeat spam reliably, b) generate a minimum of false positives, c) not be rendered obsolete by new spammer techniques and d) be deployed and managed inexpensively.

This white paper, sponsored by Abaca, discusses the scope of the spam problem, some of the techniques that spammers use to defeat conventional anti-spam technologies, and the innovative approach that Abaca has developed to stop the vast majority of spam.

The Spam Problem is Solved…Not!

A BRIEF HISTORY OF SPAM

The first record of email spam dates back as far as 1978 and, although spam began in earnest in 1994, the recent history of the spam "problem" actually began about 2002. In early 2002, spam represented about 16% of all email sent over the Internet; by early 2008, spam represents between 87% and 95% of all email.

However, the proportion of email represented by spam masks a larger problem – the absolute volume of spam that is sent on a daily basis. For example, while just a few billion spam messages were sent each day in 2002, today roughly 100 billion spam messages

traverse the Internet on a typical day. Further, spam volumes can grow rapidly over a very short period of time, such as the doubling of spam volume that occurred between May and November 2006, coupled with spam 'spikes' during the Christmas season and at other, seemingly random, times.

WHY IS SPAM SO BAD?

There are a variety of problems caused by spam:

- **Bandwidth constraints**
  Spam entering an organization's network consumes network bandwidth that could otherwise be used for legitimate purposes. As spam volumes increase, particularly as newer types of spam consume even more bandwidth on a per-message basis, bandwidth is consumed for non-legitimate purposes, in many cases requiring the deployment of larger data pipes simply to maintain acceptable performance.

- **Storage requirements**
  Similarly, as more spam enters an organization more of this content must be stored for review in spam quarantines. Given that spam is typically stored for at least 30 days for employees to review the content for false positives, increases in spam entering an organization inevitably lead to greater storage requirements.

- **Loss of employee productivity**
  While some believe that loss of employee productivity is a serious problem for most organizations, Osterman Research has found that this is actually a relatively minor problem in the overall context of the spam problem. However, it is an issue for some organizations, particularly smaller ones that do not filter spam adequately at the server or gateway.

- **Other problems**
  There are a variety of other problems related to spam, including phishing attempts to purportedly come from a valid source, such as a bank, but instead direct recipients to enter their confidential information on a phisher's Web site; some employees spending time perusing products and services offered in spam; links contained in spam messages that could direct users to harmful or offensive Web sites; and the like.

SPAMMERS ARE SMART

The keys to why the spam problem is so bad and is getting worse can be distilled down to three factors:

- Most spammers are smart
- Spammers can make money even with extraordinarily low conversion rates
- They have at their disposal lots of money for development of new delivery techniques

Given that the money being made by spammers, which totals tens of millions of dollars per year, can fund newer and better techniques for distributing spam, this provides a ready source of funding for spam development. Among the techniques that spammers use to distribute their content are:

- **Filter-circumvention techniques**
  Among the more basic techniques used by spammers are simple obfuscation techniques to trick filters into missing spam messages, such as misspelling keywords, introducing valid text like Bible verses into spam messages, using various HTML techniques to trick filters into not recognizing offensive content, Bayesian poisoning (including strings of random words to throw off spam filters) and other techniques. Two common techniques are simple misspellings of words, such as replacing an "i" with a "1"; or inserting comment tags between the letters of words in an HTML-based spam message. The latter technique, in particular, can defeat some spam filters, since the comment tags are not visible in the message displayed to humans and so the obfuscated word appears intact.

- **Botnets**
  Traditionally, spammers sent large amounts of spam from a relatively few sources that were easy to identify and block. To get around these efforts, spammers have established botnets that consist of millions of 'zombie' computers – home and corporate personal computers that are infected with a worm, virus or Trojan horse that allows them to be controlled by a spammer or other remote entity. Using botnets, which spammers can rent for spam campaigns, small numbers of messages per day can be sent from each of thousands of computers, effectively helping to keep each zombie under the radar of their respective ISPs or network administrators.

  The advantages to spammers of using botnets are that they can avoid detection by Internet Service Providers and others that look for large numbers of messages sent from individual computers, which typically indicates spam activity; and the ability to avoid blockage of their content even if large number of zombies are stopped.

  Starting in earnest in the 2006-2007 timeframe, spammers began using newer spamming techniques in an effort to defeat spam-filtering technologies.

- **Newer types of spam**
  Starting in earnest in the 2006-2007 timeframe, spammers began using newer spamming techniques in an effort to defeat spam-filtering technologies. For example:

  - **Image-based spam**
    Text is represented in one or more images often using unusual fonts, randomized backgrounds, background 'snow', slanted lines of text, fuzziness and other distortions to defeat more conventional spam-filtering technologies. Image spam is particularly bad for recipients, since each message is typically five to 10 times larger than a conventional, text-based spam message.

  - **Spam with attachments**
    Similar to image spam, but using PDF files, Excel worksheets or ZIP files as payloads to carry the spam content.

The bottom line is that spammers are clever and they are well funded. Spamming is clearly a growth industry, and newer and better techniques will be developed in an effort to defeat conventional spam-filtering technologies.

TECHNOLOGY-BASED SOLUTIONS WILL BE KEY

There have been a number of non-technological attempts to defeat spam, including the CAN-SPAM Act in the United States, various statutes enacted in a variety of nations around the world, a handful of legal actions directed against high-profile spammers, and a number of state laws that have made spamming illegal. However, none of these actions has proved effective in stopping the growing volume of spam reaching corporate networks or in thwarting the development of newer techniques designed to defeat spam filters.

Clearly, the key to stopping spam will be technology-based solutions, not legislation or legal prosecution of spammers. However, not all anti-spam technologies are created equal. Some are better than others either in spam capture efficiency and/or in generating a minimal number of false positives. While conventional spam-filtering technologies can stop a large proportion of spam, spammers continue to battle against even the cutting edge of these technologies, necessitating newer and better techniques to stop the problem.

Abaca's Solution to the Spam Problem

FOCUSES ON RECEIVER REPUTATION, NOT SENDER REPUTATION

Among the more innovative techniques currently available to defeat spam is Abaca's filtering technology. Instead of relying on the reputation of the sender or scanning the content of incoming email, the Abaca system determines the reputation of email *recipients* based on the proportion of spam that they receive.

The concept of receiver reputation is based on the fact that different people receive different amounts of spam and legitimate email. When analyzing a message, each receiver's percentage of spam versus legitimate email (his or her reputation) is an estimate of whether the message is spam or legitimate. Essentially, if the message is sent to users who typically receive a high percentage of spam, the message is more likely to be spam. However, if the message is sent to users who typically receive a low percentage of spam, the message is more likely to be legitimate. Aggregating the reputations of all recipients of a particular message, therefore, is equivalent to combining those users' rating power to estimate the legitimacy of the sender and the message. In a receiver reputation system, the key determinant of whether a message is spam or legitimate is not the identity of the sender or the content of the email, but the reputations of the email recipients, individually and collectively.

THE UNDERLYING TECHNOLOGY

The core engine behind Abaca's technology is ReceiverNet, a patent-pending, receiver reputation-based approach to detect spam. The technique is new, unique and revolutionary.

ReceiverNet is based on a sophisticated mathematical formula that uses receiver reputations to precisely differentiate spam from legitimate messages. A message is considered more likely to be legitimate if it is sent to recipients that typically receive a lower percentage of spam than the average recipient. Conversely, a message is considered more likely to be spam when sent to recipients that typically receive a higher percentage of spam than the average recipient.

It is not necessary to manage complicated rules, whitelists, or blacklists with the Abaca system. Because message ratings are based on each user's overall legitimate/spam ratio (as measured by the system), users do not need to help the system identify spam other than to express personal preferences, if they so desire. The system learns and becomes more accurate on its own by tracking the legitimate/spam statistics for each protected user. Spam detection becomes more accurate as more users join the ReceiverNet Network.

For example, if a message is sent to 100,000 protected users, the system has the rating power of 100,000 receiver reputations to rate the sender and the message. In practice, a spam attack is typically blocked before a protected user receives the first email. By the time a spammer has sent three messages, there is a 99.9 percent probability that the spam message will be blocked.

WHAT ARE THE ADVANTAGES?

The Abaca Filtering Technology offers a number of advantages compared to conventional spam-filtering technologies:

- **Very high accuracy**
  Abaca's Filtering Technology has consistently demonstrated a spam capture efficiency of 99.7%, or capture of 299 out of every 300 spam emails. More important, however, is that the technology has demonstrated an extremely low level of false positives: 0.0126%, or 126 false positives for every one million messages.

- **Spammers cannot defeat the system**
  Because spammers must send large volumes of content in order to achieve their desired return, they cannot avoid detection by the Abaca system. For example, spammers cannot obfuscate the recipient/TO field.

Abaca's Filtering Technology has consistently demonstrated a spam capture efficiency of 99.7%, or capture of 299 out of every 300 spam emails. More important, however, is that the technology has demonstrated an extremely low level of false positives.

- **Rapid response**
  ReceiverNet ratings are based on the 25 most recent emails for each sender. The result is that the system can achieve highly accurate reputation scores within just a few messages.

- **Errors are easy to find**
  The rare false positive is easily identified through examination of a daily spam report and within the quarantine, which displays spam messages by their probability ranking. Those with the lowest probability of being spam are at the top of the list, making false positive identification a simple and quick task. The spam report contains only those emails that lack sufficient statistical information to make a definitive judgment. This is typically less than 1% of the total number of emails received. So a user who gets 1,000 emails per day will typically need to review fewer than 10 messages.

- **Minimal hardware requirements**
  In the latest implementation, Abaca's Filtering Technology can process more than 22,000 messages per second using a single, 2.4GHz processor.

- **The ability to stop outgoing spam**
  The technology that allows Abaca to stop spam from entering an organization can also be used to stop outgoing spam, as in the case of a compromised, 'zombie' machine in a corporate network; or in a content-filtering use case.

- **Other advantages**
  There is no user or filter training required by the Abaca system; the system is completely language independent; and reputation scores are, for 99% of spam, either very unlikely or very likely to be spam – fewer than one percent of spam messages fall into a 'gray' area.

REAL WORLD EXAMPLES

Abaca's performance has been tested in a variety of actual use cases, as shown below:

- Community Health Partnership (CHP), based in Eau Claire, Wisconsin, provides assisted living services to the elderly, those with physical disabilities and others. Prior to deploying Abaca's Email Protection Gateway, the company's chief executive was receiving more than 400 spam messages per day and the finance department was receiving between 500 and 1,000 such messages each day – each of the 450 employees in the organization was spending at least 15 minutes per day dealing with spam. Immediately upon commencing its 30-day trial of the Abaca technology, CHP realized a 99% drop in the amount of spam it was receiving and has seen a roughly 25% increase in its bandwidth because it processes so much less spam now.

- Roggen Telephone Cooperative Company, based in rural Colorado, was receiving 30,000 to 40,000 spam messages per hour, resulting in a 95% load on its email servers despite the fact that it had an anti-spam solution in place. After deploying the Abaca Filtering Technology, the load on its email servers dropped to 0.5% on the first day after deployment.

- Cassatt, an enterprise software and services firm in San Jose, California, has 100 employees, six Microsoft Exchange servers and was receiving 3,000 spam messages on a typical day. Using competing anti-spam technologies, the best Cassatt could achieve was an 85% spam capture rate. However, after installing Abaca's system, which took less than one hour, the company has achieved a greater than 99% spam capture rate.

Summary

The spam problem is bad and is getting worse, consuming greater shares of network storage and bandwidth, sapping employee productivity and necessitating the deployment of newer and better technologies simply to maintain the status quo. Conventional spam-filtering technologies are effective to a point, but cannot keep up with the newer and more innovative technologies that spammers are developing to send their content.

Abaca has developed a new paradigm for defeating spam. Instead of relying on analysis of the content of spam or the reputation of its senders, the Abaca system uses an algorithm that analyzes the reputation of spam *recipients*. Spammers cannot defeat the system, it operates on a minimal hardware configuration, and it achieves extraordinarily high spam capture rates and very low false positives.