

The Advantages of Using Traffic-Shaping Techniques to Control Spam

An Osterman Research White Paper

Published January 2007

Sponsored by



MAILCHANNELS



Why This White Paper Will Be Worth Your Time

Spam is the scourge of anyone who uses email. While spam today represents four out of every five email messages sent, the absolute volume of spam continues to increase at a rapid pace – 2006 alone, for example, will see more than a 100% increase in the absolute volume of spam sent and received because of the use of botnets and increasing spammer sophistication.

The result of spam is not simply more junk in users' mailboxes which results in lost productivity. Spam attacks can easily bring down mail servers when sent in sufficient quantities, resulting in lost email and excessive expenditures for hardware, software and IT staff time.

While content filtering is a useful tool for eliminating spam, it is quite CPU-intensive. To combat rapidly growing volumes of spam with content filtering is expensive, since more infrastructure must be deployed to adequately manage the growing flood of spam entering individual email systems.

What is needed is a method of examining email senders before their content reaches anti-spam systems. By determining the reputation of individual senders, the amount of email that these senders are allowed to send can be increased or throttled back, resulting in a dramatic decrease of spam to mail servers.

This white paper discusses the spam problem and presents MailChannels' approach to throttling spam known as Traffic Control.

Some Basics on the Spam Problem

Although the spam problem actually began in 1994, the problem with spam really began in earnest somewhere around early 2002. At that time, spam represented about 16% of all email traffic in North America. Today, spam represents about 80% of all email.

However, the problem is actually much worse than these statistics would suggest. Because of growing volumes of both email and spam, the absolute volume of spam is dramatically greater than it used to be. For example, spam volumes approximately doubled between May and November 2006.

While spam today represents four out of every five email messages sent, the absolute volume of spam continues to increase at a rapid pace – 2006 alone, for example, will see more than a 100% increase in the absolute volume of spam sent and received because of the use of botnets and increasing spammer sophistication.

Spam volumes will continue to grow at a rapid pace with no end in sight, since the economics of spamming are decidedly with the spammers. The incremental costs associated with sending an additional 10 million or 100 million spam messages are negligible, particularly given the growth of botnets as described below, resulting in a situation that gives spammers virtually no financial incentive to target their customers in any meaningful way. The result is that spammers send enormous amounts of email, since they depend on a sale only for every 10,000 to 30,000 or more emails sent.

There have been a variety of techniques used to block spam over the past several years, including simple keyword matching, Bayesian analysis, reputation analysis and other techniques. However, the volume of spam, coupled with increasing spammer sophistication in bypassing even the best spam-filtering techniques, has created a situation in which an improved method of spam-blocking is required. In short, while content-filtering techniques will continue to be necessary tools in the fight against spam, more efficient and less CPU-intensive methods must be used to combat growing volumes of spam more efficiently.

Virtually all organizations have deployed capabilities to deal with spam. However, many of these first- and second generation systems are struggling under the rapidly growing burden of the increasing quantity and sophistication of spam. To manage the problem, organizations must either deploy more servers and the IT staff resources necessary to manage them, or they must come up with a better and more efficient method of combatting the spam problem.

The Growing Problem with Botnets

One of the most challenging aspects of dealing with spam has been the growth of 'botnets', the primary source of spam today. A botnet consists of tens of thousands of 'zombie' computers – typically home computers with a broadband connection that have become infected with malware that allows a remote party to control their behavior. Because home users are generally lax about maintaining good anti-virus and anti-spyware defenses, the botnet problem has become an epidemic of sorts, so much so that the number of zombie computers is growing by tens of thousands each day.

Botnets tax email server resources by sending very large amounts of spam in a very short period of time. This flood of

While content-filtering techniques will continue to be necessary tools in the fight against spam, more efficient and less CPU-intensive methods must be used to combat growing volumes of spam more efficiently.

email can dramatically increase the load on email servers and spam-blocking systems, resulting in service delays or even outages in some cases.

Botnets represent the best of all scenarios for spammers: 1) they can send email from literally hundreds of thousands of computers simultaneously, resulting in spam 'campaigns' that can result in literally 100 million or more spams sent within a short period of time; and 2) spammers send a relatively low volume of spam from each zombie computer, making detection quite difficult for ISPs and individual users. Blocking botnet attacks is extremely difficult because of the enormous number of computers involved, their geographic dispersion all over the globe and the fact that new zombies are created continually, making detection of zombies a somewhat futile exercise.

90% of spammers will give up their attempt to send spam to a particular address if a connection cannot be established within the first 30 seconds. Rapid connections are a necessity for spammers, since they depend on sending very large volumes of mail.

Spammers' Behavior is Bad and Typically Easy to Spot

The good news for those intent on stopping spammers is that their behavior is typically easy to spot. Spammers tend to ignore protocol rules and they conduct dictionary harvest attacks in an effort to gather new and fresh email addresses. By contrast, legitimate senders typically exhibit good behavior by complying with protocol rules, sending email that is first passed through outbound spam filters, they send email primarily to legitimate addresses, and so on.

What Works to Stop Spammers

One of the interesting findings in research about spammers is that they often are quite impatient when their SMTP connections are slowed down. For example, 90% of spammers will give up their attempt to send spam to a particular address if a connection cannot be established within the first 30 seconds. Rapid connections are a necessity for spammers, since they depend on sending very large volumes of mail.

How Traffic Control Works

MailChannels' Traffic Control has been developed as a means of controlling the flow of email, which results in dramatically reduced volumes of spam.

What is Traffic Control?


Traffic Control is an email proxy. Instead of connecting directly to a mail server, senders of email connect to Traffic Control which, in turn, connects to mail servers. One of the key advantages of Traffic Control is that it merely slows email

traffic and does not block it outright. This allows the use of a greater number of blacklists, since legitimate domains mistakenly placed on a blacklist will not see their legitimate traffic blocked as with conventional real-time blacklisting.

How Traffic Control Addresses the Problem of Botnets

The primary mechanism used by Traffic Control to manage spam is the reputation and behavior of senders. This information is used to control the quality of service that a sender receives. For example, a sender that follows SMTP protocol rules, does not send large numbers of email to invalid email addresses and is generally well behaved will establish a good reputation and receive good quality of service. Conversely, poorly behaved senders will have a low reputation score and will receive low quality of service.

MailChannels has developed a patent-pending technique called 'real-time SMTP multiplexing', which mitigates the threat from botnets.



The Impact of Traffic Shaping in the Real World

As noted above, slowing down spammers before they have a chance to deliver their email will result in a reduction of 90% of spam without incurring the risk for false positives. Because valid senders typically will not notice a delay of 30 seconds in email delivery across the Internet, the impact of the delay will be felt by spammers, not valid senders. The

result is a dramatic reduction in the amount of spam that ends up in both a spam quarantine and in users' mailboxes.

Controlling Botnets with Traffic Control

MailChannels has developed a patent-pending technique called 'real-time SMTP multiplexing', which mitigates the threat from botnets. The technique ensures that even if thousands of zombie computers are connected to Traffic Control simultaneously, only a small and manageable quantity of email traffic flows into a mail server. The result is that email bursts from botnets are eliminated, preventing spammers from bringing down email servers by flooding them with traffic.

A Service Provider Case Study

Sunflower Broadband, which serves residential customers in several communities in Kansas, manages 20,000 email accounts. Prior to the company's deployment of Traffic Control, 1.5 million email messages were processed each day, fewer than one-third of which were legitimate. Traffic spikes, most of which occurred after midnight, would cause denial-of-service attacks, resulting in system administrators being paged to correct the problem. While Sunflower was very pleased with its PureMessage anti-spam system, this solution is very CPU-intensive, since it must examine the content of each email.

To address its spam-related problems, Sunflower had planned to invest up to \$10,000 in new servers. However, the deployment of Traffic Control has resulted in the postponement of these purchases, as well as in a seven-fold increase in the number of concurrent sessions that the ISP can handle.

Summary

Spammers are aggressive and smart and continue to use new techniques to send ever increasing volumes of spam. The result is not only an enormous quantity of spam ending up in spam quarantines and users' mailboxes, but also email servers that are brought down by excessive quantities of unwanted content and increasing investments in servers and software to keep up with the growing deluge of spam.

What is needed, therefore, is a technique that blocks spam based on spammers' readily identifiable behavior without the CPU-intensive activity of content filtering. By applying a 'pre-filtering' analysis of senders' reputation, content filtering

What is needed, therefore, is a technique that blocks spam based on spammers' readily identifiable behavior without the CPU-intensive activity of content filtering. By applying a 'pre-filtering' analysis of senders' reputation, content filtering systems can operate much more efficiently, since the vast majority of spam is blocked before reaching email servers.

systems can operate much more efficiently, since the vast majority of spam is blocked before reaching email servers.

MailChannels has developed Traffic Control, a system that will categorize email senders' reputation and slow email traffic from senders who have a reputation for sending spam. However, Traffic Control does not block email, allowing real-time blacklists to be used without the risk of false positives.

© 2007 Osterman Research, Inc. All rights reserved.

No part of this document may be reproduced in any form by any means, nor may it be distributed without the permission of Osterman Research, Inc., nor may it be resold by any entity other than Osterman Research, Inc., without prior written authorization of Osterman Research, Inc.

THIS DOCUMENT IS PROVIDED "AS IS". ALL EXPRESS OR IMPLIED REPRESENTATIONS, CONDITIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OR FITNESS FOR A PARTICULAR PURPOSE, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE DETERMINED TO BE ILLEGAL.