

MailMarshal SMTP 2006

Anti-Spam Technology

August, 2006

Contents

Introduction	2
Multi-layered spam detection and management	2
SpamCensor: Marshal's unique heuristic filter	2
URLCensor: Live URL blacklist checking	3
External DNS blacklists: Query your favorite blacklist	3
CountryCensor: Analyze country of origin	3
Zero Day: For large-scale new spam outbreaks	4
DHA: Protection against Directory Harvest Attacks	4
TextCensor: Create your own custom scripts	4
And a host of other features	4
How effective is MailMarshal at blocking spam?	5
What about false positives?	5
Enabling end users to manage spam	6
Putting it together: a rules-based approach . . .	6
. . . .and a suite of management options	6
Conclusion	6

This whitepaper gives an overview of the anti-spam technology used by MailMarshal. While MailMarshal is great at blocking spam, it is much more than an anti-spam solution. MailMarshal provides organizations with the means to control all incoming and outgoing email content, including spam, viruses, text, and attachments within a rules-based framework.

Introduction

This whitepaper gives an overview of the anti-spam technology used by MailMarshal SMTP 2006. While MailMarshal is great at blocking spam, it is much more than an anti-spam solution. MailMarshal provides organizations with the means to control all incoming and outgoing email content, checking for spam, viruses, specific text and attachments within a rules-based framework.

You can envisage MailMarshal as an email toolkit, where a multitude of tools can be deployed to check and manage your email. When it comes to anti-spam, *detection* is important. But so is *management*. MailMarshal uses technologies that enable high spam detection rates with exceptionally easy and flexible administration. And it does this within the context of an integrated email content management package. MailMarshal offers organizations an effective and flexible means to control spam, with a rapid return on investment.

Multi-layered spam detection and management

Spam is constantly evolving as spammers employ ever more sophisticated techniques to evade filters. Also, the *rate* of spam evolution appears to be increasing. Spam is now the domain of sophisticated money-generating organizations with a high level of technical prowess. The result is that no single piece of anti-spam technology is likely to be effective against spam. Your best approach is to adopt a multi-faceted solution. MailMarshal has an array of anti-spam tools that together, layer-by-layer, identify and deal with spam.

MailMarshal Multi-Layered Anti-Spam Protection



SpamCensor: Marshal's unique heuristic filter

The SpamCensor is at the core of the MailMarshal anti-spam solution. The SpamCensor is an heuristic filter that consists of approximately 3000 individual tests for spam. It is a scoring-based system where rules work in combination to end up with a total score of the 'spamminess' of a message. Once the score reaches a threshold, action is taken like quarantining or tagging the message.

The SpamCensor is built and automatically updated on a weekly basis by members of Marshal's TRACE (Threat Research and Content Engineering) team. This is a group of security analysts who examine live email streams for patterns and typical spam "traits".

WHITEPAPER – MailMarshal SMTP 2006 – Anti-Spam Technology

The experience and knowledge of the TRACE team is a key factor in the effectiveness of the SpamCensor. This knowledge is backed up by a number of proprietary tools and systems that automatically “machine-learn” patterns and assign scores to rules to optimize performance. The SpamCensor is self-contained, stand-alone and requires no work by an administrator. If you used nothing else in MailMarshal, the SpamCensor alone would result in very good spam detection.

Sometimes customers ask why the SpamCensor is not updated as often as other anti-spam solutions or their virus scanner. The simple answer is that it is not a signature-based system. In such systems, each signature is an independent entity and equates to a single message. In the SpamCensor, rules are *heuristic* and *interdependent*. Individual signature updates are unnecessary. A weekly release cycle is used to carefully craft, train and test each SpamCensor prior to release. This approach results in a highly *predictive* filter that is very good at detecting tomorrow’s spam, as well as today’s.

URLCensor: Live URL blacklist checking

One very effective way of blocking spam is to extract domain information from URLs in the body of the message, and check that information against one or more of the URL blacklists available. This method is known as SURBL (Spam URL Realtime Blacklisting). MailMarshal’s implementation of URL blacklisting is called URCensor, and it contains the following advanced features:

- Ability to use any number of URL blacklist databases.
- Can also check IP addresses of extracted domains.
- Local caching of results, to increase speed and to reduce the load on DNS servers.
- Handles obfuscated URLs.
- Ability to use multiple URL blacklists in a policy-based framework (i.e. only block a message if it is listed on two or more URL blacklists).

External DNS blacklists: Query your favorite blacklist

Another established way of blocking spam is to check the sender against a blacklist of known spamming hosts. There are quite a number of blacklists of varying quality and availability. As the services use DNS as the method of querying their servers, they are also often referred to as DNS blacklists.

The effectiveness of this method of blocking spam is entirely dependent on the quality of your chosen list, and how often it is updated. The services sometimes attract criticism because, occasionally, legitimate email hosts can find themselves unwittingly on the list, and it becomes difficult to send email to them. Even so, a good DNS blacklist can be effective at blocking spam and should form part of an overall anti-spam strategy.

MailMarshal has integrated DNS blacklist support and ships with some already configured. Administrators can apply one or more DNS blacklists as needed.

CountryCensor: Analyze country of origin

The CountryCensor is unique technology developed by Marshal which allows a mail administrator to identify the countries through which a message has traveled, and handle it accordingly. The technology analyzes IP addresses to identify the country of origin. Internet authorities allocate blocks of IP Addresses to each country, and CountryCensor uses a database of these allocations to identify the country of origin. You can then set policy in MailMarshal to scrutinize messages from some countries more closely, or deny messages

WHITEPAPER – MailMarshal SMTP 2006 – Anti-Spam Technology

from certain countries altogether. For example, if you never do business with anyone in the Republic of Zamunda you can block any messages originating from that country. You can also grant individual exceptions such as the domains of any customers you might have in that country.

Zero Day: For large-scale new spam outbreaks

From time to time, large-scale or worldwide spam outbreaks occur. Marshal has the ability to push out targeted Zero Day protection measures. If you enable the Zero Day policy, these measures are automatically applied. Zero Day spam measures are designed to provide interim protection until the SpamCensor is next updated.

DHA: Protection against Directory Harvest Attacks

DHA prevention guards your system against Directory Harvest Attacks (DHAs). Spammers use DHAs to determine valid email addresses at your domain. This mechanism can detect a DHA, drop the connection from the connecting server and blacklist the server for a specified length of time. .

TextCensor: Create your own custom scripts

MailMarshal has an integrated text scanner called the TextCensor. These scripts can be used in numerous ways to stop spam – they provide you with the flexibility to look for almost anything. The TextCensor engine has advanced capability, and its scripts support:

- Boolean operators e.g. AND, OR, NOT
- Proximity operators e.g. NEAR, FOLLOWEDBY, INSTANCES
- Phrase weightings e.g. "FREE!" might be given a higher weighting than "buy now"
- Targeting of different parts of message. The scripts can be limited to searches the header, body, subject lines, or attachments

And a host of other features.....

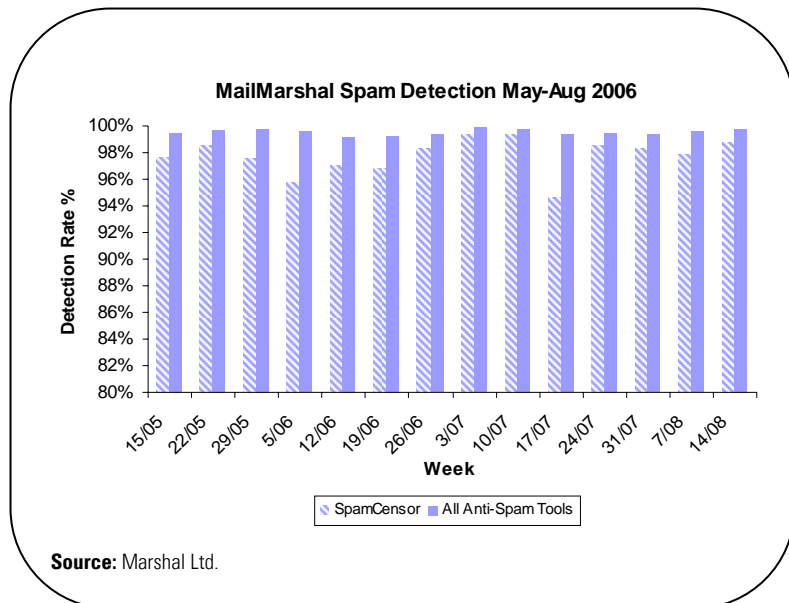
The features listed above represent the core of MailMarshal's anti-spam capabilities. But MailMarshal also uses a range of other methods to deal with spam as well:

- User groups for managing whitelists and blacklists
- Anti-spoofing
- Anti-relay
- Denial of Service (DOS) protection
- Receiver HELO connection rules
- Reverse DNS lookups
- Regular expression header matching and rewriting
- Advanced custom category scripts for combining regular expression matching and TextCensor scripts.

When all these technologies are combined together into one layered solution, MailMarshal is one of the most effective and easy to manage solutions for defeating spam on the market.

How effective is MailMarshal at blocking spam?

The Marshal TRACE team closely monitors the performance of anti-spam on live MailMarshal servers subjected to streams of incoming spam. The following chart illustrates its performance:



There are a few points to note here:

- The *average* detection rate of SpamCensor alone during this period was *97.8%*.
- The SpamCensor combined with the other anti-spam tools (e.g. URLEnsure and DNS Blacklists), MailMarshal's effectiveness against spam increases to *99.5%*.
- Overall effectiveness remains constant at *~99.5%*. In weeks where the SpamCensor rate drops, the other tools act as a buffer to keep the overall detection rate up.

What about false positives?

It is relatively easy to block spam. The trick is to do it with a minimum of false positives. MailMarshal achieves spam false positive rates of better than 0.01% false positive rate (1 in 10,000 messages). This is an excellent figure when benchmarked against other best-in-class anti-spam systems. An important point to note here is that the 0.01% false positive rate is based on *all* email, including subscription bulk email. Legitimate, person-to-person business email has an even better rate approaching 0.0001% or 1 in 1,000,000 messages.

The critical issue with false positives is managing them. Regardless of what solution you have, at some point a false positive will occur. MailMarshal makes it easy to manage spam and ensure that a legitimate email is not lost.

Enabling end users to manage spam

No solution is completely foolproof with regard to false positives. The key thing is managing them and ensuring that these messages are not lost or deleted. This is one of the areas where MailMarshal has significant advantages. Quarantined spam can be subjected to the end user Spam Quarantine Management (SQM). This system periodically sends a digest email to each user with a list of blocked spam that was addressed to them. The user can then link directly from the email to the online SQM system and manage their own spam as they wish – releasing quarantined messages and defining safe senders or senders they want to block in the future.

Putting it together: a rules-based approach....

Flexibility is one of MailMarshal's strengths. You can harness the power of MailMarshal's pre-configured spam technology, such as the SpamCensor and get great results. However, MailMarshal also provides the ability to customize rules for every site. Administrators can combine rule elements together to create a policy that is greater than the sum of its parts.

Here is a simple example: Spam messages are often not very large – almost invariably less than 125Kb in size. This fact can be used in conjunction with TextCensor scripts and whitelists to create an accurate rule, as in the following example from the MailMarshal rule wizard:

```
When a message arrives
Where message is incoming
  Except where addressed either to or from 'Excluded Users'
  Except where addressed from 'Friendly ListServers'
Where message size is less than '125 Kb'
And where message is categorized as 'Spam'
Move the message to 'Spam'
```

....and a suite of management options

Once a message has been determined as spam, administrators must decide what action to take with it. MailMarshal provides a wide array of possible actions for the maximum flexibility. Here is a sample of the many options available:

- Move the message to a quarantine folder
- Copy the message to a folder
- Send a notification message
- Write a custom log message to the database for later reporting
- Rewrite the message headers – e.g. mark the subject line with [SPAM]
- Route the message to another host
- Pass the message to another rule for processing
- Delete the message

Messages can be monitored and controlled using the MailMarshal Console. Administrators can have any number of consoles and they can be configured to control only specified user groups if desired.

Conclusion

MailMarshal provides email administrators with the technology for controlling spam. It combines leading, new technology with traditional anti-spam approaches. Above all, as stressed in this paper, it provides anti-spam capability in a highly flexible and easy to use solution. In short, MailMarshal has:

- The range and depth of technology for maximum *accuracy* in spam detection

WHITEPAPER – MailMarshal SMTP 2006 – Anti-Spam Technology

- A rules-based approach for maximum *usability*
- A suite of management options for maximum *flexibility* in the enterprise.

Marshal is committed to providing the best possible email content security solution. Our research and development team is working continuously to improve MailMarshal - adding greater accuracy of detection and more intuitive and flexible management functionality. For more information on MailMarshal please contact your Marshal reseller or sales representative, or visit www.marshal.com.

THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT ARE FURNISHED UNDER AND ARE SUBJECT TO THE TERMS OF A LICENSE AGREEMENT OR A NON-DISCLOSURE AGREEMENT. EXCEPT AS EXPRESSLY SET FORTH IN SUCH LICENSE AGREEMENT OR NON-DISCLOSURE AGREEMENT, MARSHAL LIMITED PROVIDES THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SOME JURISDICTIONS DO NOT ALLOW DISCLAIMERS OF EXPRESS OR IMPLIED WARRANTIES IN CERTAIN TRANSACTIONS; THEREFORE, THIS STATEMENT MAY NOT APPLY TO YOU.

This document and the software described in this document may not be lent, sold, or given away without the prior written permission of Marshal, except as otherwise permitted by law. Except as expressly set forth in such license agreement or non-disclosure agreement, no part of this document or the software described in this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, or otherwise, without the prior written consent of Marshal. Some companies, names, and data in this document are used for illustration purposes and may not represent real companies, individuals, or data. This document could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein. These changes may be incorporated in new editions of this document. Marshal may make improvements in or changes to the software described in this document at any time.

© 2006 Marshal Limited, all rights reserved.

U.S. Government Restricted Rights: The software and the documentation are commercial computer software and documentation developed at private expense. Use, duplication, or disclosure by the U.S. Government is subject to the terms of the Marshal standard commercial license for the software, and where applicable, the restrictions set forth in the Rights in Technical Data and Computer Software clauses and any successor rules or regulations.

Marshal, MailMarshal, the Marshal logo, WebMarshal, Security Reporting Center and Firewall Suite are trademarks or registered trademarks of Marshal Limited or its subsidiaries in the United Kingdom and other jurisdictions. All other company and product names mentioned are used only for identification purposes and may be trademarks or registered trademarks of their respective companies.



Marshal's Worldwide and EMEA HQ
Marshal Limited,
Renaissance 2200,
Basing View,
Basingstoke,
Hampshire RG21 4EQ
United Kingdom

Phone: +44 (0) 1256 848080
Fax: +44 (0) 1256 848060

Email: emea.sales@marshal.com

Americas
Marshal Inc.
5909 Peachtree Dunwoody Road, NE,
Suite 770,
Atlanta,
GA 30328
USA

Phone: +1 404 564-5800
Fax: +1 404 564-5801

Email: americas.sales@marshal.com
www.marshal.com | info@marshal.com

Asia-Pacific
Marshal Software (NZ) Ltd
Suite 1, Level 1, Building C
Millennium Centre
600 Great South Road
Greenlane, Auckland
New Zealand

Phone: +64 9 984 5700
Fax: +64 9 984 5720

Email: apac.sales@marshal.com