

Marshal TRACE Report

2006 Year in Review

Contents

Introduction	2
The Launch of TRACE Center	2
Spam	3
Phishing	8
Malware	9
How did the Anti-Spam Industry Fare in 2006?	10
Spam’s Impact on Marshal Customers in 2006	10
What Are the Spam Fighting Lessons From 2006?	11
Conclusion	12

This report has been prepared by the Marshal Threat Research and Content Engineering Team (TRACE). It is a review of the trends and changes in global spam during the course of 2006.

The intended audience for this report is interested parties in the wider Marshal community (customers, partners, media and analysts).

This report addresses the major spam related changes in 2006. Why these changes happened, and how Marshal succeeded in protecting its customers against spam in 2006.

Introduction

2006 witnessed some of the most significant developments in the world of spam in recent years. At the World Economic Forum in January 2004, Bill Gates famously predicted that spam would be “a thing of the past” by the year 2006.

Hopes were high that new anti-spam technology such as Bayesian logic and reputation-based filtering would end spam for good. Others thought that developing legislation against spamming would help. During 2005 many in the anti-spam community felt they were getting on top of the problem.

However, in 2006 the opposite happened. Spam bounced back and today is more abundant than ever. Why did this happen and how did the world’s most powerful technology mogul get his prediction so wrong? What are the lessons to be learnt and how does all this affect Marshal customers?

This report seeks to answer some of those questions.

The Launch of TRACE Center

After separating from NetIQ in December 2005, Marshal set about re-establishing its software development resources. As part of this effort, the Marshal TRACE Team was formed, tasked with researching and developing technologies to combat Internet-based threats.

The TRACE Team researches the areas of spam, phishing, and general malware. It is responsible for the producing and releasing SpamCensor updates for MailMarshal, and for feeding the results of research into ongoing product development.

During the year, the Team built up a substantial infrastructure and volume of data to support its research. In October 2006, the TRACE Center Website was launched.

Data and analysis from the TRACE Team can be viewed at any time at www.marshal.com/trace.


Spam

The year 2006 marked some huge changes in the spam landscape.

Spam Volume Rises Markedly

The overall volume of spam increased in 2006 by some 200% as spammers evolved and found success with new methods and tactics.

The first part of year was relatively normal until a huge surge in spam volume occurred in mid-October. You can see this trend in the following chart of Marshal's Spam Volume Index – measuring the volume of spam received by a sample of honeypot domains.



Marshal Spam Volume Index

This marked increase in spam was directly attributable to the successful distribution and deployment of new malware, notably the Stration/Warezov email worm, and the SpamThru Trojan, among others.

In addition to increased numbers of spam messages, spam got larger (owing to image spam) causing up to a threefold increase on the volume of email traffic companies had to handle. The sheer volume of new spam caught some companies and ISPs short as they struggled to upgrade their systems to cope.

Image Spam

The other major change in 2006 was the explosion of image spam, which seeks to display the spammer's message in an attached image. Though the idea itself is not new, spammers came up with creative new ways of using images and this type of spam began to be sent on a scale not seen before. The technique is designed to beat traditional keyword anti-spam technologies by making the spam message unreadable to a machine, but still readable by a human.

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review

Image spam increased from 15% (of all spam) at the beginning of 2006 to peak at 50% in late December.



Percentage of Image Spam

Image spam rapidly evolved during 2006 as spammers experimented with new techniques and variations in an attempt to defeat anti-spam technologies – with some success. Some anti-spam products struggled to maintain detection rates.

In the beginning images were clean and simple, like this stock spam example:

Golden Apple Oil and Gas, Inc. ([GAPJ](#))

THIS STOCK IS EXTREMELY UNDERVALUED!

GAPJ - is our **NEXT HOT PICK**, which we feel is most undervalued stock we have ever featured and should out perform all other picks.

Golden Apple Oil and Gas, Inc. ([GAPJ](#)) *Oil and Gas sector

Current Price: **\$1.15**

Short Term Price Target: **\$3.50**

Long Term Price Target: **\$7.00**

Status: **Strong Buy**

*300+% profit potential short term

RECENT HOT NEWS released MUST READ ACT NOW

VPHOENIX, Jan. 13, 2006 (PRIMEZONE) -- Golden Apple Oil and Gas, Inc. ([GAPJPK - News](#)) is pleased to announce that it has secured the first \$100,000.00 placement as they finalize Arrangements for additional financings. The company is planning additional announcements Regarding Directors, Advisors and Management in the next few days

* ^ * ^ *

About Golden Apple Oil and Gas, Inc.:

Headquartered in Phoenix, Arizona, Golden Apple is an independent oil and gas producer With a focus on North American properties. The Company applies advanced technologies To systematically explore and develop its oil and natural gas opportunities.

Watch this One Trade Tomorrow! GO GAPJ!

An early 'clean' example of image spam.

By year end, the images had changed markedly. We observed extreme, over-the-top methods used to confuse and defeat image analysis technology and checksum-based anti-spam solutions. These methods included randomization with dots, alternating lines and background colors as well as "captcha" methods where words are bent and blurred so that they are less recognizable by OCR technology (Optical Character Recognition).

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review




A later example of "extreme" image spam showing random dots, colors, shapes and also uneven text designed to fool OCR technology.

We also saw complex "puzzle spam" where the image was broken into pieces, making it harder to analyze through image recognition technologies. The arrangements of images were automatically re-assembled by the user's email client; as far as the user was concerned, they were viewing one image. A variation on this method that we called "ransom note" spam made each letter in the spam message was an individual image, often in a different font. This gave the message the appearance of a ransom note assembled from different letters cut out of a newspaper.

Composition and Origin of Spam


The most common category of spam was Health. This is spam promoting pharmaceuticals such as weight loss pills and performance enhancement drugs. Health spam normally accounts for around 40%-45% of spam. The second biggest category is Product spam which advertises everything from replica Rolex watches to copier toner and cheap software. Product spam normally accounts for 15%-20% of all spam.




Spam by Type 2006

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review

In terms of source by country, the USA continued its hold on first place. However, China, Poland, Spain and South Korea all vied for the second tier places in 2006 as significant sources of spam – demonstrating that many computers in those countries are compromised spam “zombies”. By Continent, Europe emerged as the dominant source of all spam in 2006, significantly ahead of the next closest region, Asia.



Spam by Country 2006



Spam by Continent 2006

Stock Spam

The stock spam “pump and dump” phenomenon continued in 2006. Stock spam promotes a company’s stock, passing itself off as a “hot” stock tip. It has appeared in force only in the last two years. Today stock spam represents around 15% of all spam.

Stock spam doesn’t ask you to directly buy a product or service. Rather it takes the form of bogus “advice” on the prospects of a particular company, along with price quotes and recommendations. These hot tips advertise genuine companies listed on stock exchanges.

This “stock hype by spam” often causes price spikes in the stock, allowing the spammers and early purchasers of the stock to cash in for a profit. Unfortunately the price usually drops substantially within days and unsuspecting or naive people inevitably get burnt.

The worrying thing is that some people, unassociated with the original perpetrators, are obviously acting on stock spam in an organized way. This exacerbates the problem, and

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review

allows the spammers to drop any pretense making their stock spam look like “legitimate” advice.

The possibility that spammers can influence financial markets in this way is of wider concern, especially as there is little evidence to suggest that this phenomenon is waning. Bigger scams involving bigger companies and more credible information might just be around the corner.

URL Links in Spam Decline

An interesting trend in 2006 was the decline of URL links in spam. A few years ago almost all spam contained an http web site link so the user could click and find out more about the product advertised. In 2004, the percentage of spam with a clickable URL link was 96%. The ubiquitous nature of URLs in spam gave rise to the useful anti-spam practice of URL blacklisting and such services as SURBL (Spam URI Realtime Blocklist). Anti-spam products such as MailMarshal could simply extract the URL from the spam message and query the blacklist database.

However, the rapid rise in stock spam and image spam has affected URLs in spam. Stock spam is unique in that it is not advertising a product as such and therefore spammers have little need to provide a URL. And image spam provides a means for spammers to hide their URLs from anti-spam devices by providing the URL in the image and giving instructions to the user to type it into their browser.

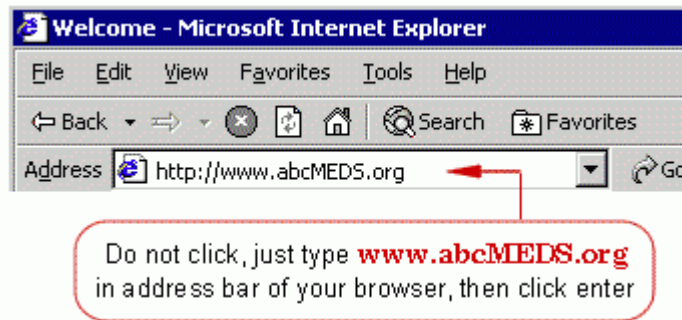


Image spam with URL instructions instead of an actual link.

In 2006 the proportion of spam with clickable URLs dropped markedly, and now sits at around 55%.

Percentage of Spam with URLs	
2004	96%
2005	86%
Jan '06	86%
Jun '06	73%
Jan '07	55%

URL blacklisting is now not as powerful as it once was. The situation is yet another example of how spammers constantly morph their practices to overcome anti-spam technology.

Legislation proves ineffective


Anti-spam legislation does not seem to have helped stem the flow of spam. Between 2003 and 2006 many countries announced new anti-spam legislation. The most prominent of these new laws was the U.S. “Can Spam Act” which led the way in 2003. Other Western

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review

countries such as Australia, Japan and the United Kingdom quickly followed suit, and by the end of 2006 many countries had legislation in place. Unfortunately, legislation has so far been ineffective, resulting in nothing more than a tiny number of symbolic prosecutions, and no visible reduction in spam. Today, spammers seem almost completely unworried and unaffected by all attempts at controlling the problem through legislation.

Phishing

Phishing also evolved considerably during 2006. Overall, phishing volumes remained constant (as a percentage of spam) at around the 1% level for the year.



Phishing, Percentage of all Spam

What we did notice however, was that phishing become slicker, more sophisticated, and more targeted.

Typical phishing emails were quick to use image spam to get around text scanning, and to look highly professional.

Banks Are Still Major Targets

Targets for phishing continued to revolve around major banks and of course companies like eBay and PayPal. However, during 2006, smaller banks in smaller countries were targeted. Instead of blasting the globe, phishers increasingly concentrated on individual regions and countries, presumably looking for fresh targets.

The intention of phishing messages is to induce recipients into unwittingly revealing personal information about themselves (usually their login name and password).

In the case of banks, the objective is to steal the victim's money straight out of their bank account. In the case of companies like eBay, the intention was to hijack other people's accounts and use them to advertise non-existent goods or take advanced deposits for goods and leave the real user to deal with the police and the ripped-off buyer.

However, in the end, banks were the favorite target, with US bank Fifth Third being the target of over 30% of all phishing attacks in the last half of 2006.

Mules

Associated with phishing were attempts to recruit 'mules'. There are the seemingly random job offers spammed to your inbox offering easy money with minimal work.

The offers are from criminals seeking to launder their ill-gotten gains. While the email phishers grab the limelight for stealing personal data, these scammers are in the background putting the stolen IDs to use.

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review

The scam aims to convert stolen personal and financial data into cash, and is often quite elaborate with real looking companies and websites.

Mules help to keep goods flowing through a distribution system, and they insulate the real criminals from the police by making it harder to track financial transactions.

Malware

The people behind malware grew increasingly professional in 2006. Malware is no longer the realm of 'script kiddies' seeking fame, or seeing how much damage they can do. Rather it is now squarely the domain of professionals seeking to steal sensitive information, or to use PCs as members of their botnets. In some cases, it is clear malware authors do not use the stolen data themselves, but merely on-sell it to third parties.

This increasing professionalism coincided with more stealth. As money making operations, malware authors now want their creations to remain hidden from view and undetected. Rootkits and other stealth methods evolved and were deployed more widely in 2006.

Traditional mass mailing viruses waned in 2006, at least until the Stration/Warezov worm appeared in October. This was a true mass mailing worm with a twist – the ability to vary itself by downloading new components from the web. The anti-virus vendors struggled to keep up with the many hundreds of variants of this worm and we saw many examples slip through this layer of defense. A similar situation occurred in early 2007 with the Storm Worm. Despite this situation, many MailMarshal administrators had little problem as blocking executables at the email gateway was a simple, yet highly effective, way of stopping these worms.

Malware, Botnets and Spam

The links between malware and spam are clear and well established. Today the bulk of spam is distributed via botnets – networks of malware-infected computers controlled by the bad guys known as 'bot-herders'. Recent estimates suggest as many as 10% of all computers connected to the Internet are infected and are part of a botnet.

The basic technique is to distribute malware, entice users to execute it, and then get it to download spam sending software. The flood of spam in 2006 after October was in large part attributable to the Stration/Warezov worm. Six hours after this worm executes, it downloads further components from a remote server and duly proceeds to send spam, unknown to the owner of the PC.

Another good example is the SpamThru (or SpamBot) Trojan, associated with stock spam. Once executed, SpamThru communicates with other members of the botnet using its own custom peer-to-peer protocol. It receives instructions from a central command server and downloads spam templates, email addresses, random From: names and random phrases for hash-busting. And to top it all off, SpamThru uses a pirated copy of Kaspersky Anti-Virus to scan for and remove other malware! The sophistication of SpamThru rivals that of commercial software, and is surely a precursor of what is to come in the near future.

The spam botnets have become so successful that the traditional method of sending spam via open-relay servers has declined markedly. In December the open relay blacklist provider ORDB.org announced it was closing, noting that owing to the change in spammer tactics, open relay blacklists are no longer an effective way of preventing spam.

Microsoft Office Targeted

During 2006 numerous vulnerabilities and exploits targeted Microsoft Office. Of the 103 critical updates released by Microsoft in 2006, 40% of them were related to Microsoft Office software, notably Word, Excel and PowerPoint.

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review

Also notable was the appearance of highly targeted attacks – five or six people in a single company targeted, with Office attachments appearing to come from other people in the company.

Office appears to be an obvious and easy target. Being normally business related, firewalls and content filters tend to pass Office documents through. Microsoft Office is part of normal business life and opening them is natural – especially if they appear to come from someone you know. We are likely to see more of these Microsoft Office exploits in the near future.

How did the Anti-Spam Industry Fare in 2006?

For several years now the fight between spammers and anti-spam vendors has been an ongoing battle of cat and mouse. As soon as one side gains an advantage, the other side tries to outmaneuver that advantage or change direction.

The challenges in 2006 in the form of increased volume, image spam and hash-busting techniques caused many headaches for ISPs and companies everywhere. By the end of 2006 the percentage of spam in inbound mail reached all time highs of 80%-90% of all email. Corporate and ISP email systems simply drowned in spam.

Some anti-spam vendors struggled to maintain detection rates during 2006. Anti-spam filters based around a single technology were exposed. Checksum or signature-based products, or products that were primarily text analyzers, performed poorly when the advanced image spam took off. Furthermore, the increasing effectiveness of botnets has turned normally trustworthy PCs into sources of spam, reducing the effectiveness of reputation-based spam filters.

In response to the image spam problem, some vendors developed complex image analysis methods (OCR or Optical Character Recognition) to attempt to read the words inside the spam images. While seemingly a good idea, this technique did little for accuracy and created performance bottlenecks owing to the resource intensive nature of the technology.


In summary, the spammers arguably won the battle in 2006 or, at the very least, scored a few points.

Spam's Impact on Marshal Customers in 2006

Despite what was happening elsewhere, MailMarshal SMTP performed extremely well against all forms of spam and image spam in particular.

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review

During 2006 MailMarshal's anti-spam catch rate steadily improved and by year's end averaged 99.5%. The key reasons for this performance were the research input from the TRACE team into the SpamCensor and the multi-layered nature of MailMarshal's approach.



MailMarshal Spam Detection 2006

The SpamCensor increasingly targeted spambot 'traits' or 'footprints' in 2006, improving its ability to block similar spam types over and over again. The content of different spam campaigns may differ, but other traits and patterns do not.

In addition to this, MailMarshal's anti-spam processing speed remained high and messages were quickly identified and classified as spam without bottlenecks occurring in message delivery.

We received very few complaints from our customers regarding spam in 2006, and in many cases minor configuration and environmental issues were at fault – all easily rectified.

Inevitably, some spam got through and we did hear about this from some customers. In one case, an email administrator received a complaint from the CEO that he had 5 spam messages in his inbox. In response, the administrator ran a report and showed the CEO the 1,500 spam messages that MailMarshal had blocked for him that day.

In another case, the CEO of a major banking organization initially wanted to throw out MailMarshal and replace it with another solution. The CEO changed his opinion after seeing a spam blocking report from MailMarshal and speaking to a leading industry analyst about what sort of spam capture rate his company should be striving for. He ended up calling and congratulating us.

What Are the Spam Fighting Lessons From 2006?

So what lessons did we learn from 2006?

Lesson 1: *A multi-faceted, multi-layered approach to spam filtering is critical. Not only can it increase accuracy, it provides different ways to detect spam making the system more resilient and less susceptible to outbreaks of new types of spam.*

As mentioned previously, spammers employ a range of methods to defeat spam filters. If the spammers employ new techniques to defeat a particular technology, it is

WHITEPAPER – Marshal TRACE Report – 2006 Year in Review

essential to have a fallback. MailMarshal performed well against all forms of spam in 2006 because it is a layered, multi-faceted solution. Customers should use all the tools available to them, including SpamCensor, multiple DNS blacklists, URLCensor, and other technology.

Lesson 2: Optimize your system for the fastest email processing.

The huge increase in spam volume now requires sorting through dozens of pieces of junk to find the one or two legitimate email messages that you actually want to receive. Filtering must be fast and efficient, and not bottleneck email delivery time. Delays in delivering legitimate email heavily marginalize the benefits that spam filtering is intended to provide in the first place. To be successful against spam it is not enough to be accurate, you need to be fast too.

Customers should optimize their system by rejecting spam up-front (see Lesson 3), organizing rules in a logical manner, and optimizing DNS Blacklist queries for minimum lag.

Lesson 3: The more spam that you can reject up-front at the SMTP level the better. Being able to reject spam at the Receiver is a key advantage of MailMarshal that many customers do not fully utilize.

MailMarshal has the ability to apply a range of checks at the SMTP level, prior to the body of the message being downloaded. Using Receiver Rules, MailMarshal can apply DNS Blacklists, and check whether the connecting SMTP mail server uses industry standards to identify itself. Receiver blocking has the distinct advantage of saving bandwidth, processing time and disk space.

Lesson 4: If the system is missing significant amounts of spam, get the configuration checked prior to throwing the product away.

The vast majority of 'missed spam' cases Marshal dealt with in 2006 were configuration issues, notably over-zealous sender whitelists. In most cases small configuration changes are all that is necessary to easily rectify the problem.

Lesson 5: Adapt or die. Systems must adapt to changing spammer tactics.

The pace of change in spammer tactics in 2006 was astonishing. Will this continue? You bet! Spam filtering systems must themselves be flexible, and able to adapt. MailMarshal was able to adapt throughout 2006 to changing spammer tactics. Product changes were delivered through weekly updates to the SpamCensor and normal product releases.

Conclusion

The year 2006 might be remembered as 'the year spam came back'. Spammers got serious and introduced new tactics and fast changing technology. The spammers arguably gained some degree of success in terms of defeating anti-spam technology. The level of sophistication went up a notch as new malware was released which created more spambots which in turn pushed spam volumes to record highs.

Unfortunately, spam and the professional malware organizations behind it show no signs of letting up. It's our challenge to keep ahead of the game. Marshal is committed to vigilantly maintaining and improving spam detection. We aim to introduce a number of new spam fighting technologies and services in 2007.

We hope that you have found this report interesting and informative. If you have any questions or comments we would very much like to hear them. You can email us at trace@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.

© 2007 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
info@marshal.com | www.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