SEPTEMBER 2014

virus BULLETIN

Covering the global threat landscape

VBSPAM COMPARATIVE REVIEW SEPTEMBER 2014 – SUMMARY

INTRODUCTION

In this short version of the September 2014 VBSpam report, we provide a summary of the results of the 33rd VBSpam test as well as some information on 'the state of spam'. The main point of note from the test results is that, with one exception, products performed very well.

We also look at a possible correlation between the number of URLs present in the body of an email and the likelihood of emails being blocked. Somewhat surprisingly, while we would expect more URLs to mean that the email is more likely to be blocked, we didn't find this to be the case.

THE VBSPAM TESTS

The VBSpam tests started in May 2009 and have been running every two months since then. They use a number of live email streams – of which the spam feeds are provided by *Project Honey Pot* and *Abusix* – which are sent to participating solutions in parallel to measure their ability to block spam and to correctly identify various kinds of legitimate emails. Products that combine a high spam catch rate with a low false positive rate (the percentage of legitimate emails that are blocked) achieve a VBSpam award, while those doing this exceptionally well earn a VBSpam+ award.

This month's VBSpam test saw 16 anti-spam solutions and two DNS blacklists on the test bench. Filtering more than 140,000 emails over a 16-day period, all but one full solution performed well enough to achieve a VBSpam award[1] – and seven of them achieved a VBSpam+ award. Once again, these results demonstrate that, while spam remains a problem that cannot be ignored, there are many solutions that do a very good job of mitigating it.

[1] Given that DNS blacklists are supposed to be included in an anti-spam solution rather than run on their own, it is not reasonable to expect such products to meet our strict thresholds. Thus, while the DNS blacklist solutions included in the test did not achieve a VBSpam award, they certainly didn't 'fail' the test.

THE RESULTS

Despite the fact that one solution failed to achieve a VBSpam award, the overall performance of the products on test this month was once again good. Excluding the single outlier, the average spam catch rate increased – more than making up the drop in catch rate reported in July. Again excluding the outlier, the average participating product missed fewer than one in 500 emails in our spam streams.

It should be noted that these numbers don't necessarily translate to performance in a live environment – where there will be a number of other factors that affect filters' performance and, just as importantly, their perceived performance. What matters is that a 99.90% catch rate is better than a 99.70% one and, in an average live environment, a product that achieves a 99.90% catch rate in our tests is likely to perform better, even if the actual catch rates differ.

There was a small increase in the overall false positive rate seen in this test, but nothing to worry about. Meanwhile, the performance on newsletters (a feed that includes anything from weekly consumer offers to daily news digests) improved – though for obvious reasons, distinguishing legitimate commercial emails from illegitimate ones remains one of the trickiest parts of running a spam filter.

Among the 15 products that passed the test, seven – *Bitdefender*, *ESET*, *GFI*, *Libra Esva*, *Mailshell*, *Netmail Secure* and *OnlyMyEmail* – achieved a VBSpam+ award for blocking more than 99.5% of spam, while blocking no legitimate emails, and very few newsletters.

OnlyMyEmail once again achieved the highest spam catch rate (the hosted solution missed only one out of 133,020 spam emails), closely followed by *Libra Esva* and *Mailshell*.

ISSN 1749-7027

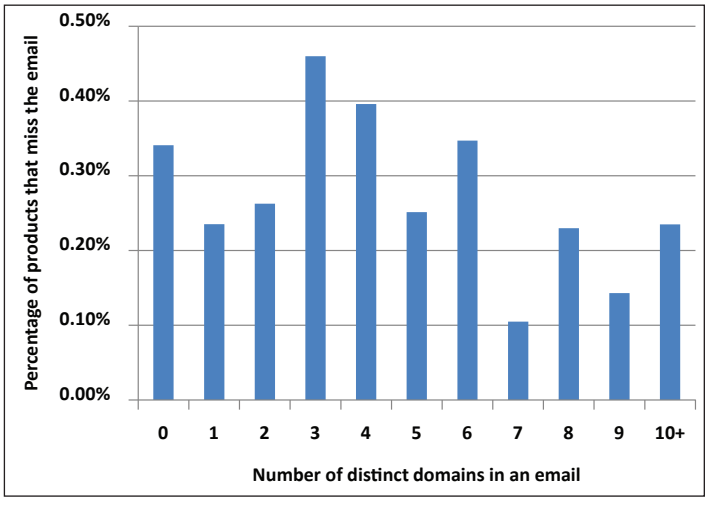

Figure 1: Does the number of domains present in an email correlate to the likelihood of the email being blocked?

Two products, *Mailshell* and *ESET*, deserve a special mention as the only products that managed not to clock up any false positives on either the ham feed or the newsletter feed.

URLS IN EMAILS

Most spam filters use a number of techniques to detect spam, and keeping track of spammy or malicious URLs is one of them. For practical reasons, this is mainly done by keeping track of malicious domains.

Spammers have used many techniques to avoid their emails being blocked this way – for instance by using compromised legitimate domains, URL-shorteners (such as bit.ly or goo.gl), or even by not including URLs in their emails at all.

This led us to ask the question: does the number of domains present in an email correlate to the likelihood of the email being blocked?

In the graph shown in Figure 1, the horizontal axis shows spam emails containing a certain number of distinct domains[2], while the vertical axis shows the percentage of products missing the average email of this kind. Somewhat surprisingly, there isn't a significant decrease (or even a decrease at all) in this chance as the number of domains in an email increases.

Of course, with catch rates as high as they are, one wouldn't expect to find a strong correlation, and there may also have been spam campaigns with no or few URLs that were easy to block for a different reason, thus blurring our statistics. Hence, while this result is certainly interesting, it doesn't mean that spammers can safely add more URLs to their spam without increasing the chances of their emails being blocked.

[2] In line with what is common practice in URL blacklists, we only looked at the 'smallest relevant part' of a domain name, e.g. google.co.uk or google.com. So www.google.com and mail.google.com were considered to be the same domain.

TABLE AND GRAPH

Note that in the table on page 3, products are ranked by their 'final score'. This score combines the spam catch rate, false positive rate and newsletter false positive rate in a single metric. However, readers are encouraged to consult the in-depth report for the full details and, if deemed appropriate, use their own formulas to compare products.

In the VBSpam quadrant, the products' spam catch rates are set against their 'weighted false positive rates', the latter being a combination of the two false positive rates, with extra weight on the ham feed. An ideal product would be placed in the top right corner of the quadrant.

Editor: Martijn Grooten
Chief of Operations: John Hawes
Security Test Engineers: Scott James, Tony Oliveira, Adrian Luca
Sales Executive: Allison Sketchley
Editorial Assistant: Helen Martin
Perl Developer: Tom Gracey
Consultant Technical Editors: Dr Morton Swimmer, Ian Whalley

Tel: +44 (0)1235 555139. Fax: +44 (0)1865 543153
Email: editorial@virusbtn.com
Web: http://www.virusbtn.com/

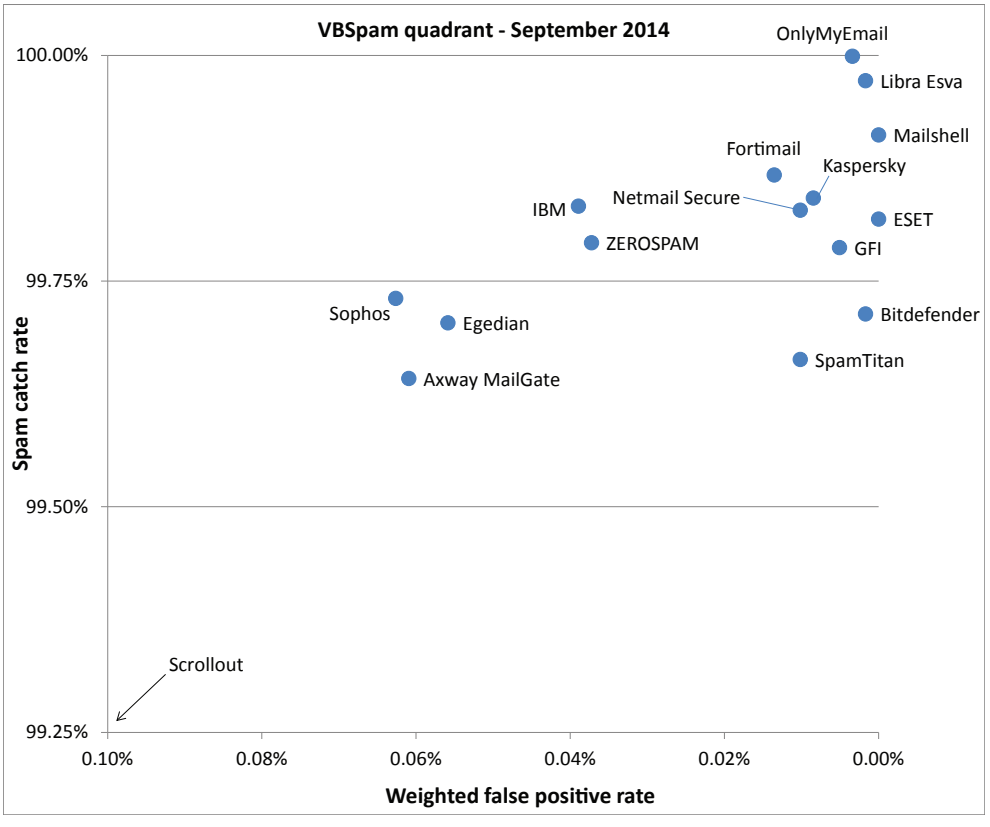

Product name	True negatives	False positives	FP rate	False negatives	True positives	SC rate	Final score
OnlyMyEmail	11748	0	0.00%	1	133019	99.999%	99.98
Libra Esva	11748	0	0.00%	37	132983	99.97%	99.96
Mailshell	11748	0	0.00%	117	132903	99.91%	99.91
ESET	11748	0	0.00%	241	132779	99.82%	99.82
FortiMail	11747	1	0.01%	176	132844	99.87%	99.80
Kaspersky LMS	11747	1	0.01%	210	132810	99.84%	99.80
Netmail Secure	11748	0	0.00%	228	132792	99.83%	99.78
GFI	11748	0	0.00%	283	132737	99.79%	99.76
Bitdefender	11748	0	0.00%	381	132639	99.71%	99.71
IBM	11744	4	0.03%	222	132798	99.83%	99.64
SpamTitan	11747	1	0.01%	448	132572	99.66%	99.61
ZEROSPAM	11744	4	0.03%	276	132744	99.79%	99.61
Egedian	11741	7	0.06%	394	132626	99.70%	99.42
Sophos	11741	7	0.06%	358	132662	99.73%	99.42
Axway	11743	5	0.04%	476	132544	99.64%	99.34
Scrollout	11737	11	0.09%	1793	131227	98.65%	97.73
Spamhaus ZEN*	11748	0	0.00%	14042	118978	89.44%	89.44
Spamhaus DBL*	11747	1	0.01%	89645	43375	32.61%	32.57

**The Spamhaus products are partial solutions and their performance should not be compared with that of other products. Please refer to full report for full product names and details.*