SEPTEMBER 2014

virus BULLETIN

Covering the global threat landscape

VBSPAM COMPARATIVE REVIEW SEPTEMBER 2014

INTRODUCTION

As testers, we have an immense advantage compared with those managing the email security products we are testing.

We are able to decide that we can't vouch for the test conditions being optimal during certain time periods and thus exclude the corresponding emails from the test. We can even turn off the feed for some time, to allow us to perform urgent maintenance. We have made use of both options on several occasions in the past.

Email is designed to work even when a mail server is temporarily unavailable, but spam filter vendors would lose customers if their products didn't work 24/7. They would lose customers even more quickly if they told those customers that they couldn't vouch for the accuracy of filtering 'between Wednesday afternoon and Thursday morning'.

Still, despite having this relative advantage, we always aim for a smooth test and were thus pleased that, for the first time in several months, and thanks to some improvements to our test set-up made over the summer, we neither needed to extend the test period, nor had to exclude large chunks of email from the corpus on this occasion.

There was some delay in the publication of this report though, largely due to the team taking time out of their regular schedule to run and attend the Virus Bulletin conference at the end of September. As always, there were a number of presentations at the conference both on and relating to spam – which, if anything, renewed our enthusiasm for running these tests.

This month saw 16 full solutions on the test bench, all but one of which achieved a VBSpam award. There were seven solutions that didn't block a single legitimate email and, combining this with a high spam catch rate and a low newsletter-blocking rate, earned a VBSpam+ award.

THE TEST SET-UP

The VBSpam test methodology can be found at http://www.virusbtn.com/vbspam/methodology/. As usual, emails were sent to the products in parallel and in real time, and products were given the option to block email pre-DATA (that is, based on the SMTP envelope and before the actual email was sent). Three products chose to make use of this option.

For those products running on our equipment, we use *Dell PowerEdge* machines. As different products have different hardware requirements – not to mention those running on their own hardware, or those running in the cloud – there is little point comparing the memory, processing power or hardware the products were provided with; we followed the developers' requirements and note that the amount of email we receive is representative of that received by a small organization.

To compare the products, we calculate a 'final score', which is defined as the spam catch (SC) rate minus five times the weighted false positive (WFP) rate. The WFP rate is defined as the false positive rate of the ham and newsletter corpora taken together, with emails from the latter corpus having a weight of 0.2:

WFP rate = (#false positives + 0.2 * min(#newsletter false positives , 0.2 * #newsletters)) / (#ham + 0.2 * #newsletters)

Products earn VBSpam certification if the value of the final score is at least 98:

SC - (5 x WFP) ≥ 98

Meanwhile, products that combine a spam catch rate of 99.5% or higher with no false positives and no more than 2.5% false positives among the newsletters earn a VBSpam+ award.

ISSN 1749-7027

THE EMAIL CORPUS

The test period started at 12am on Saturday 16 August and ran for 16 consecutive days, ending at 12am on Monday 1 September.

The test corpus consisted of 145,109 emails. 133,020 of these emails were spam, 50,100 of which were provided by *Project Honey Pot*, with the remaining 82,920 emails provided by *spamfeed.me*, a product from *Abusix*. They were all relayed in real time, as were the 11,748 legitimate emails ('ham') and 341 newsletters.

Figure 1 shows the catch rate of all full solutions throughout the test. To avoid the average being skewed by poorly performing products, the highest and lowest catch rates have been excluded for each hour.

The exclusion of the weakest performance from the average matters this time, because of one outlier. Hence, while the overall average catch rate was slightly lower on this occasion, most products actually saw their catch rates increase – which the chart also shows.

Only once did the (corrected) average drop below 99% – this was caused by a few instances of German language recruitment spam. These and similar emails were among the hardest to filter in this test, but it was an instance of Bosnian political spam that took the title of the most difficult to filter spam email in this test.

RESULTS

In the text that follows, unless otherwise specified, 'ham' or 'legitimate email' refers to email in the ham corpus – which excludes the newsletters – and a 'false positive' is a message in that corpus that has been erroneously marked by a product as spam.

Because the size of the newsletter corpus is significantly smaller than that of the ham corpus, a missed newsletter has a much greater effect on the newsletter false positive rate than a missed legitimate email has on the false positive rate for the ham corpus (e.g. one missed email in the ham corpus results in an FP rate of less than 0.01%, while one missed email in the newsletter corpus results in an FP rate of almost 0.3%).

Axway MailGate 5.3.1

SC rate: 99.64%

FP rate: 0.04%

Final score: 99.34

Project Honey Pot SC rate: 99.34%

Abusix SC rate: 99.83%

Newsletters FP rate: 3.2%

vb VERIFIED SPAM virusbtn.com

Axway is gradually climbing the VBSpam rankings, with the *MailGate* virtual appliance yet again increasing its spam catch rate this month. Among the fewer than 500 spam emails that the product did miss, we

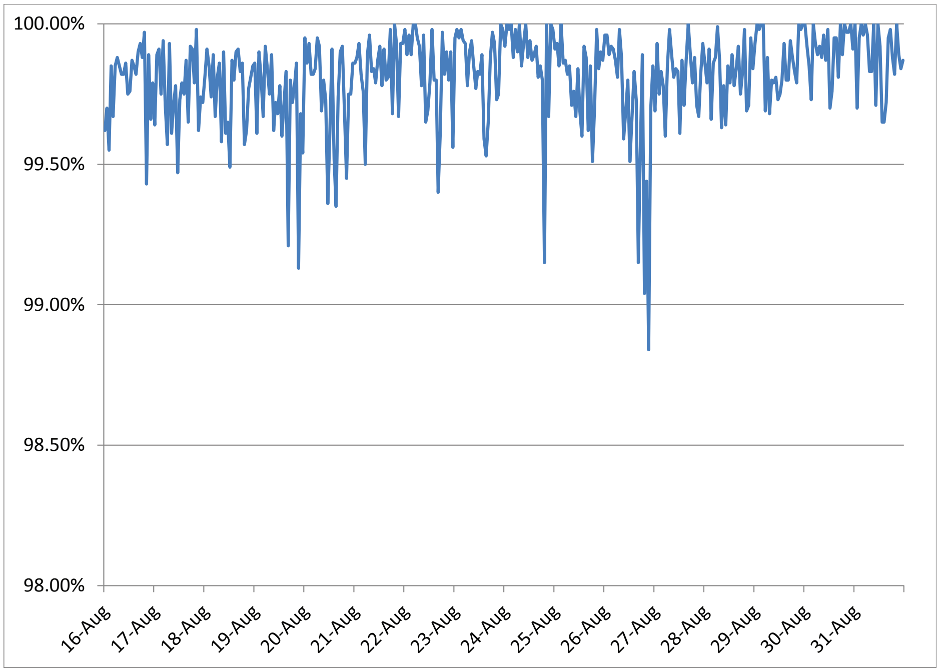

Figure 1: Spam catch rate of all full solutions throughout the test period.

noticed quite a few written in Japanese, as well a significant number of English-language dating spam emails.

At the same time, both the product's false positive rate and its newsletter false positive rate decreased slightly, resulting in an increased final score and a well-deserved VBSpam award – the product's fourth.

Bitdefender Security for Mail Servers 3.1.2

vb VERIFIED SPAM + virusbtn.com

SC rate: 99.71%

FP rate: 0.00%

Final score: 99.71

Project Honey Pot SC rate: 99.31%

Abusix SC rate: 99.96%

Newsletters FP rate: 0.3%

A fair amount of spam in Asian character sets caused *Bitdefender*'s spam catch rate to drop a little this month. It would be unfair to call this bad news though, as the catch rate remained high, and there were yet again no false positives for the Romanian product.

Thus, not only does *Bitdefender* achieve its 33rd VBSpam award in a row, it also continues the unbroken run of VBSpam+ awards that it started in January 2013.

Egedian Mail Security

vb VERIFIED SPAM virusbtn.com

SC rate: 99.70%

FP rate: 0.05%

Final score: 99.42

Project Honey Pot SC rate: 99.43%

Abusix SC rate: 99.87%

Newsletters FP rate: 0.9%

Among the spam emails missed by *Egedian Mail Security* were several in a campaign that spread links to adult content on compromised websites. Few other products had difficulty with these emails, but this is merely an indication of the fact that each product is different: when it came to its spam catch rate, *Egedian*, like other products, saw an improved performance this month.

Thus, despite a slight increase in the false positive rate (but a sharp drop in that of the newsletters), the final score for the virtual solution from *Profil Technology* increased a little further and the product earned its third VBSpam award.

ESET Mail Security for Microsoft Exchange Server

vb VERIFIED SPAM + virusbtn.com

SC rate: 99.82%

FP rate: 0.00%

Final score: 99.82

Project Honey Pot SC rate: 99.73%

Abusix SC rate: 99.87%

Newsletters FP rate: 0.0%

Japanese dating spam, English adult content spam, and malicious spam in German were among just 243 emails missed by *ESET*. This false negative rate alone meant that the product performed better than average, but in fact, the 243 missed emails were all that we could find wrong with *ESET*'s performance – the product didn't miss any legitimate emails, or any emails in the hard-to-filter newsletter corpus.

This clean sheet for the product earns it its sixth VBSpam+ award.

Fortinet FortiMail

vb VERIFIED SPAM virusbtn.com

SC rate: 99.87%

FP rate: 0.01%

Final score: 99.80

Project Honey Pot SC rate: 99.71%

Abusix SC rate: 99.96%

Newsletters FP rate: 0.9%

We recently tidied the server rack that houses all of the VBSpam machines, and it was nice to see the *FortiMail* appliance we set up well over five years ago still humming along nicely. This month sees the product finish its 100000th test, and although that number is written in binary, it is still an impressive achievement – in each of the 32 tests *Fortinet*'s appliance has taken part in, it has achieved a VBSpam award.

There was no VBSpam+ award in the bag this time, due to a single false positive, but with a better-than-average performance in all metrics, *Fortinet*'s developers have every reason to celebrate.

GFI MailEssentials

vb VERIFIED SPAM + virusbtn.com

SC rate: 99.79%

FP rate: 0.00%

Final score: 99.76

Project Honey Pot SC rate: 99.66%

Abusix SC rate: 99.87%

Newsletters FP rate: 0.9%

An eclectic mix of 283 spam emails, as well as three newsletters from the same sender, were all that *GFI MailEssentials* misclassified in this test, meaning that, as in the last 18 tests in which it has participated, the *Windows* solution achieves a VBSpam award.

Product name	True negatives	False positives	FP rate	False negatives	True positives	SC rate	Final score
Axway	11743	5	0.04%	476	132544	99.64%	99.34
Bitdefender	11748	0	0.00%	381	132639	99.71%	99.71
Egedian	11741	7	0.05%	394	132626	99.70%	99.42
ESET	11748	0	0.00%	241	132779	99.82%	99.82
FortiMail	11747	1	0.01%	176	132844	99.87%	99.80
GFI	11748	0	0.00%	283	132737	99.79%	99.76
IBM	11744	4	0.03%	222	132798	99.83%	99.64
Kaspersky LMS	11747	1	0.01%	210	132810	99.84%	99.80
Libra Esva	11748	0	0.00%	37	132983	99.97%	99.96
Mailshell	11748	0	0.00%	117	132903	99.91%	99.91
Netmail Secure	11748	0	0.00%	228	132792	99.83%	99.78
OnlyMyEmail	11748	0	0.00%	1	133019	99.999%	99.98
Scrollout	11737	11	0.09%	1793	131227	98.65%	97.73
Sophos	11741	7	0.06%	358	132662	99.73%	99.42
SpamTitan	11747	1	0.01%	448	132572	99.66%	99.61
ZEROSPAM	11744	4	0.03%	276	132744	99.79%	99.61
Spamhaus DBL*	11747	1	0.01%	89645	43375	32.61%	32.57
Spamhaus ZEN*	11748	0	0.00%	14042	118978	89.44%	89.44

**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.*

Moreover, and perhaps more importantly, the lack of false positives combined with a high spam catch rate and low newsletter FP rate means that *GFI* achieves its fourth VBSpam+ award.

IBM Lotus Protector for Mail Security

SC rate: 99.83%

FP rate: 0.03%

Final score: 99.64

Project Honey Pot SC rate: 99.57%

Abusix SC rate: 99.99%

Newsletters FP rate: 0.9%

Looking at the 222 spam emails *IBM Lotus Protector* missed, we see a mix of emails in English and Chinese, as well as a fair number of very short emails that give spam filters very little to base their blocking decisions on. It will be useful for the developers to know that almost all missed emails stemmed from the *Project Honey Pot* feed.

The product also erroneously blocked four legitimate emails and three newsletters – all of which were written in English – but the product's final score was more than adequate to earn the product its 17th VBSpam award and show that *IBM* continues to do a good job.

Kaspersky Security 8 for Linux Mail Server

SC rate: 99.84%

FP rate: 0.01%

Final score: 99.80

Project Honey Pot SC rate: 99.71%

	Newsletters		Project Honey Pot		Abusix		pre-DATA[‡]		STDev[†]
	False positives	FP rate	False negatives	SC rate	False negatives	SC rate	False negatives	SC rate	
Axway	11	3.2%	332	99.34%	144	99.83%			0.67
Bitdefender	1	0.3%	345	99.31%	36	99.96%			0.69
Egedian	3	0.9%	284	99.43%	110	99.87%			0.42
ESET	0	0.0%	134	99.73%	107	99.87%			0.39
FortiMail	3	0.9%	144	99.71%	32	99.96%			0.26
GFI	3	0.9%	172	99.66%	111	99.87%			0.31
IBM	3	0.9%	213	99.57%	9	99.99%			0.28
Kaspersky LMS	0	0.0%	145	99.71%	65	99.92%			0.35
Libra Esva	1	0.3%	33	99.93%	4	100.00%	120115	90.30%	0.12
Mailshell	0	0.0%	68	99.86%	49	99.94%			0.22
Netmail Secure	6	1.8%	171	99.66%	57	99.93%	119381	89.75%	0.31
OnlyMyEmail	2	0.6%	1	99.998%	0	100.00%			0.01
Scrollout	54	15.8%	118	99.76%	1675	97.98%			2.11
Sophos	2	0.6%	327	99.35%	31	99.96%			0.46
SpamTitan	1	0.3%	348	99.31%	100	99.88%			0.41
ZEROSPAM	2	0.6%	238	99.52%	38	99.95%	105009	78.94%	0.37
Spamhaus DBL[*]	0	0.0%	27531	45.05%	62114	25.09%			12.97
Spamhaus ZEN[*]	0	0.0%	12093	75.86%	1949	97.65%			4.09

[*] Spamhaus is a partial solution and its performance is not to be compared with that of other products.

[‡] pre-DATA filtering was optional and was applied on the full corpus. All *ZEROSPAM*'s false positives occurred pre-DATA; others were post-DATA.

[†] The standard deviation of a product is calculated using the set of its hourly spam catch rates.

(Please refer to the text for full product names.)

Kaspersky Security 8 for Linux Mail Server contd.

Abusix SC rate: 99.92%

Newsletters FP rate: 0.0%

We don't usually know why a product blocks a certain email, but in the case of the single false positive for *Kaspersky*'s *Linux Mail Server* product, it may well be due to a somewhat unusual URL shortener used in the body of the email. While they are useful for making URLs look prettier, shorter and easier to copy, URL shorteners are also popular among spammers as they hide the final destination of links contained within emails.

It was this single false positive – the product didn't even block any newsletters – that stood in the way of the product earning another VBSpam+ award. Still, with performance in all areas better than average, the product's developers have plenty to be pleased with.

Libra Esva 3.3.2.0

SC rate: 99.97%

FP rate: 0.00%

Final score: 99.96

Project Honey Pot SC rate: 99.93%

Abusix SC rate: 100.00%

Pre-DATA SC rate: 90.30%

Newsletters FP rate: 0.3%

Hosted solutions	Anti-malware	IPv6	DKIM	SPF	DMARC	Multiple MX-records	Multiple locations
Mailshell	Optional		√	√		√	√
OnlyMyEmail	Proprietary (optional)		√	√	*	√	√
ZEROSPAM	ClamAV			√		√	√

* OnlyMyEmail verifies DMARC status but doesn't provide feedback at the moment.

(Please refer to the text for full product names.)

Local solutions	Anti-malware	IPv6	DKIM	SPF	DMARC	Interface			
						CLI	GUI	Web GUI	API
Axway MailGate	Kaspersky, McAfee	√	√	√				√	
Bitdefender	Bitdefender	√				√		√	√
Egedian	Bitdefender	√				√		√	√
ESET	ESET Threatsense					√	√		
FortiMail	Fortinet	√	√	√		√		√	
GFI	Five anti-virus engines	√		√				√	
IBM	Sophos; IBM Remote Malware Detection			√		√		√	
Kaspersky LMS	Kaspersky	√		√		√		√	
Libra Esva	ClamAV; others optional		√	√		√		√	
Netmail Secure	Proprietary	√	√	√		√		√	
Scrollout	ClamAV			√		√		√	
Sophos	Sophos							√	
SpamTitan	Kaspersky; ClamAV	√	√	√		√		√	√

(Please refer to the text for full product names.)

Libra Esva missed just 37 spam emails. That's fewer than all but one other solution – and almost all of those 37 emails were ones that other products had difficulty with too.

Apart from these, the product misclassified a single newsletter (an offer from a Thai newspaper), but once again there were no false positives in the ham corpus for the Italian virtual solution. *Libra Esva*'s developers thus have reason to celebrate as the product earns its tenth VBSpam+ award.

Mailshell Mail Agent

SC rate: 99.91%

FP rate: 0.00%

Final score: 99.91

Project Honey Pot SC rate: 99.86%

Abusix SC rate: 99.94%

Newsletters FP rate: 0.0%

vb
VERIFIED
SPAM +
virusbtn.com

A year after its last participation, *Mailshell Mail Agent* returns to the VBSpam test – and what a return it is! The product missed just 117 spam emails – a score bettered by only two other products – and did not misclassify any emails in the ham corpus or in the newsletter corpus.

This impressive performance earns *Mailshell* a very well deserved VBSpam+ award.

Netmail Secure

SC rate: 99.83%

FP rate: 0.00%

Final score: 99.78

Project Honey Pot SC rate: 99.66%

Abusix SC rate: 99.93%

Pre-DATA SC rate: 89.75%

Newsletters FP rate: 1.8%

vb VERIFIED SPAM + virusbtn.com

A campaign of Japanese spam emails sent during the early days of the test proved tricky for *Netmail Secure*, as did various emails in a number of Western languages. This was nothing to worry about though: the virtual appliance missed fewer than one in 500 emails in the spam feed.

Moreover, while there were six newsletters that were incorrectly blocked, there were no false positives among the more than 11,700 emails in the ham corpus. As such, a VBSpam+ award is sent *Netmail*'s way – the product's sixth.

OnlyMyEmail's Corporate MX-Defender

SC rate: 99.999%

FP rate: 0.00%

Final score: 99.98

Project Honey Pot SC rate: 99.998%

Abusix SC rate: 100.00%

Newsletters FP rate: 0.6%

vb VERIFIED SPAM + virusbtn.com

Over the years, we've become accustomed to *OnlyMyEmail*'s excellent performance, but misclassifying only three emails in a corpus of more than 140,000 is still something that amazes us. Just one spam email – from a French lottery – and two newsletters (one of which, interestingly, was also in French) received the incorrect classification.

Thus, once again, *OnlyMyEmail*'s hosted solution has the highest catch rate, the highest final score and, of course, another VBSpam+ award – the product's seventh.

Scrollout F1

SC rate: 98.65%

FP rate: 0.09%

Final score: 97.73

Project Honey Pot SC rate: 99.76%

Abusix SC rate: 97.98%

Newsletters FP rate: 15.8%

This wasn't a good month for *Scrollout F1*. The open-source product missed more spam than any other full solution (though, on the *Project Honey Pot* feed, its performance was actually better than average) and had more false positives in both corpora than any other product.

This meant that the product's final score dropped below the certification threshold of 98. This may be good enough for some – and the product's past performance has certainly shown that it's capable of performing well – but for us it was reason to deny it a VBSpam award on this occasion, the product's first fail since this time last year.

Sophos Email Appliance

SC rate: 99.73%

FP rate: 0.06%

Final score: 99.42

Project Honey Pot SC rate: 99.35%

Abusix SC rate: 99.96%

Newsletters FP rate: 0.6%

vb VERIFIED SPAM virusbtn.com

A difficulty in parsing the headers added by our MTA saw a significantly lower catch rate for *Sophos*'s *Email Appliance* in the last test than its developers thought it should have. This test confirms that last month's score was indeed an anomaly, as the percentage of missed spam halved.

Unfortunately, seven false positives – emails from Russian *Linux* users and American birdwatchers – caused the product's final score to decrease a little. Still, the performance remained good and it's now up to the developers to work towards another VBSpam+ award in the next test. For now, the product's 26th VBSpam award will have to suffice.

SpamTitan 6.00

SC rate: 99.66%

FP rate: 0.01%

Final score: 99.61

Project Honey Pot SC rate: 99.31%

Abusix SC rate: 99.88%

Newsletters FP rate: 0.3%

vb VERIFIED SPAM virusbtn.com

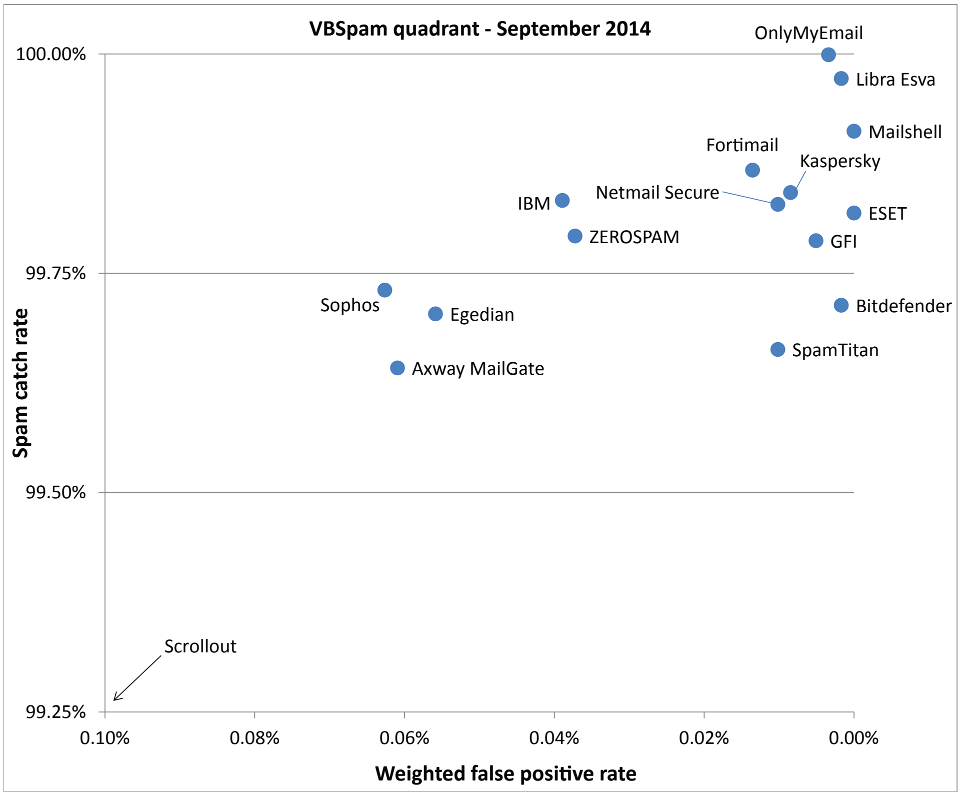

This was a good month for *SpamTitan*: after we described the product's results in the last test as 'a bit disappointing', the virtual appliance bounced back and performed well on all counts. The product missed only around one in 300 spam emails – among which were some offering dog treats and anti-aging products.

The product earns its 30th VBSpam award and, with a single false positive, it falls just a whisker short of a VBSpam+ award.

ZEROSPAM

SC rate: 99.79%
FP rate: 0.03%
Final score: 99.61
Project Honey Pot SC rate: 99.52%
Abusix SC rate: 99.95%
Pre-DATA SC rate: 78.94%
Newsletters FP rate: 0.6%

vb VERIFIED SPAM virusbtn.com

In the last test, we mentioned that *ZEROSPAM* had missed more image spam emails than any other product. Image spam isn't necessarily any worse than other kinds of spam, and the remark was meant to highlight an area in which the developers of the hosted solution could work on improving its performance a little. Perhaps they did – as the product's performance on image spam certainly saw some improvement.

Not only that, but *ZEROSPAM* slashed more than one third off its false negative rate. This was pleasing to see but, unlike the previous test, there were two false positives, meaning that the product fell short of earning a VBSpam+ award this time. Nevertheless, it earns its 16th VBSpam award (a crucial number for those of us working in computers!).

Spamhaus DBL

SC rate: 32.61%
FP rate: 0.01%
Final score: 32.57
Project Honey Pot SC rate: 45.05%
Abusix SC rate: 25.09%
Newsletters FP rate: 0.0%

Spamhaus ZEN

SC rate: 89.44%
FP rate: 0.00%
Final score: 89.44

Spamhaus ZEN contd.

Project Honey Pot SC rate: 75.86%
Abusix SC rate: 97.65%
Newsletters FP rate: 0.0%

The blacklisting of a free hosting domain meant that *Spamhaus*'s domain-based blacklist *DBL* didn't finish the test free from false positives. This demonstrates that using domain-based blacklists isn't entirely risk-free, and it emphasizes the difficulty of maintaining such a list. Indeed, the same domain was present in several spam emails in this month's corpus.

We should, of course, also note that there was an increase of over four percentage points in the percentage of emails that included a *DBL*-listed domain. At the same time, the catch rate of *Spamhaus ZEN*, a combination of various IP-based blacklists, remained just below 90%. For a solution that only looks at connecting IP addresses and that is supposed to be used in combination with other spam-filtering techniques, that remains a decent score.

CONCLUSION

For 15 of the 16 participating full solutions, this month's test results contained good news – although, knowing that most developers aim far higher than a standard VBSpam award, for many it also contains areas for improvement.

For us as testers, the best news was the fact that the test ran so smoothly. Having performed some much needed maintenance on our systems – more of which has taken place since the test has finished – we should now be able to add some new features to the test. More on those in future reports.

The next VBSpam test will run in October 2014, with the results scheduled for publication in November. Developers interested in submitting their products should email martijn.grooten@virusbtn.com.

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

© 2014 Virus Bulletin Ltd, The Pentagon, Abingdon Science Park, Abingdon, Oxfordshire OX14 3YP, England.
Tel: +44 (0)1235 555139. Fax: +44 (0)1865 543153
Email: editorial@virusbtn.com
Web: http://www.virusbtn.com/