NOVEMBER 2012

virus BULLETIN

Fighting malware and spam

NOVEMBER 2012 VBSPAM COMPARATIVE REVIEW

INTRODUCTION

All spam is bad. But some spam messages are more harmful than others.

One of the spam filters that has participated in the VBSpam tests displays a counter which indicates how much money you have saved by running the filter, on the basis that receiving one spam message means a cost of $0.04. This seems like a reasonable estimate: distraction of the user, if only for a few seconds, together with a tiny increase in the need for disk storage and network capacity, leads to a small cost.

Spam is largely a quantitative problem. No company goes bankrupt through receiving a few dating scams, and a few dozen pharmacy spam emails won't affect the balance sheet in any significant way. But because spam is sent in such vast quantities, installing a spam filter isn't merely a sound business decision, it is essential in order for an organization to continue to receive email.

However, in some cases, even a single message can have serious financial consequences. An extreme example of this is what security firm *RSA* experienced last year: an employee opened a malicious attachment that led to outsiders gaining access to the company's network – the incident is reported to have cost the company around $66 million. (It is important to point out that the email *was* actually blocked by a spam filter, but subsequently retrieved from quarantine.)

Less extreme examples of single spam messages having a financial impact on the recipient are more common: emails that claim to contain news of an important update from your bank, *PayPal* or *Twitter*, and which provide a link to what appears to be the bank/*PayPal*/*Twitter* to confirm your details but which instead leads to a phishing site. In many cases, these emails are sent in much smaller quantities than other types of spam and much more effort is made to ensure that the emails appear genuine, making them a lot harder to detect.

It was with this in mind that we decided to add a new feed to the anti-spam tests – one which specializes in phishing emails. The results from this feed show that products do indeed have more difficulty with identifying these kinds of emails than with the average spam message.

Of course, the test also included the usual ham and spam feeds, based on which 19 full solutions achieved 15 VBSpam awards and a record number of four VBSpam+ awards.

THE TEST SET-UP

The VBSpam test methodology can be found at http://www.virusbtn.com/vbspam/methodology/. As usual, email was sent to the products in parallel and in real time, and products were given the option to block email pre-DATA. Five products chose to make use of this option.

As in previous tests, the products that needed to be installed on a server were installed on a *Dell PowerEdge R200*, with a 3.0GHz dual core processor and 4GB of RAM. The *Linux* products ran on *CentOS 6.3* (*Bitdefender*) or *Ubuntu 11* (*Kaspersky*); the *Windows Server* products ran on either the 2003 or the 2008 version, depending on which was recommended by the vendor.

To compare the products, we calculate a 'final score', which is defined as the spam catch (SC) rate minus five times the false positive (FP) rate. Products earn VBSpam certification if this value is at least 97:

$$SC - (5 \times FP) \geq 97$$

Meanwhile, those products that combine a spam catch rate of 99.50% or higher with a lack of false positives earn a VBSpam+ award.

ISSN 1749-7027

As has been highlighted before, there is no objective justification for using the weight of 5 in the calculation of the final score: the spam catch and false positive rates are two distinct metrics of a spam filter and any way in which they are combined into a one-dimensional metric is arbitrary.

We use the weight of 5 to highlight the importance of false positives, without false positives becoming the single metric that determines whether products pass or fail. We have received suggestions both to increase and to decrease the weight. Readers who prefer to use a different weight – or a different formula altogether – are encouraged to do so using the numbers presented in the tables.

THE EMAIL CORPUS

As usual, the test ran for 16 consecutive days, from 12am GMT on Saturday 20 October 2012 until 12am GMT on Monday 5 November 2012.

The corpus contained 104,972 emails, 92,166 of which were part of the spam corpus. Of these, 74,021 were provided by *Project Honey Pot*, and 18,145 were provided by *Spamfeed.me*, a product from *Abusix*. They were all relayed in real time, as were the remaining emails, consisting of exactly 12,000 legitimate emails ('ham'), 202 legitimate newsletters and 604 emails from the *Wombat* corpus, more on which below.

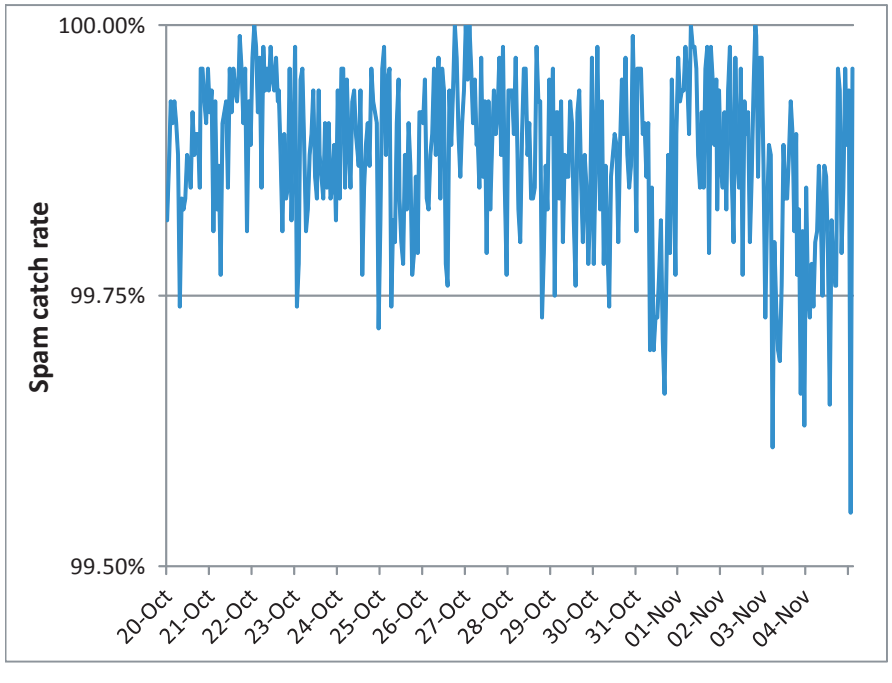

Figure 1: Spam catch rate of all complete solutions throughout the test period.

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.

Comparing the graph with those produced in previous tests, it is immediately clear that overall performance was better on this occasion than it has been in recent tests. In fact, this is the first time that the average performance has stayed above 99.5% throughout the test – a welcome change after we reported a general decline in product performance during the first half of the year.

One thing we should point out is that all spam used in our tests has been sent to spam traps. This means that, for obvious reasons, the more targeted type of spam is less likely to appear in the feeds. This kind of spam tends to be harder to filter: when the 'quantity' of a spam campaign decreases, its 'quality' (i.e. its ability to evade filters) generally increases. It is thus worth keeping in mind that performance in the VBSpam tests should only be considered in the context of the tests, and catch rates should be seen as relative to other products' catch rates rather than an accurate estimation of how much spam a product blocks in the wild.

This is why the *Wombat* feed introduced in this test is an important addition.

The feed is provided by *Wombat Security Technologies*, a vendor of anti-phishing training and filtering solutions, which was founded in 2008 by faculty members of Carnegie Mellon University – who have done a significant amount of research on phishing over the past decade.

The *Wombat* feed consists of live, run-of-the-mill phishing emails. It is skewed towards consumer-oriented phishing emails written in English, such as emails claiming to come from financial institutions or popular online services such as *Facebook* and *Twitter*. The emails in the feed generally urge the user to click a link to reverse an account suspension, confirm an update or read a new message. Of course, none of the links actually lead to a genuine legitimate site: in most cases they direct to a fake site phishing for account information, though in some cases they lead to malicious sites which install malware on the recipient's machine via a drive-by download.

While the consumer-oriented phishing emails in the *Wombat* feed are not

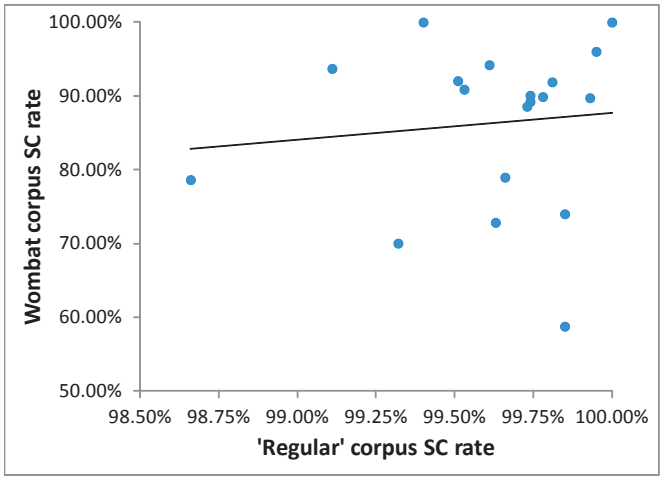

Figure 2: Correlation between performance on Wombat and 'regular' spam feeds.

representative of some of the more targeted spear-phishing emails (which tend to be harder to detect), they do reveal significant differences in performance between the products.

Unsurprisingly, for all but two products, performance on the *Wombat* feed was significantly lower than on the normal spam feed. How much lower, however, varied greatly and Figure 2 shows there is only limited correlation between products' performance on the two feeds.

For technical reasons, the emails in the *Wombat* feed were predominantly sent during the second half of the test. Participants had received feedback on their products' performance on this feed, but this happened only once and shortly before the start of the test; in some cases participants may argue that they would have performed better with one or more configuration changes. Future tests will determine whether this is indeed the case.

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.

The term 'spam' exclusively refers to those messages in the *Project Honey Pot* and *Abusix* corpora and, unless otherwise specified, the 'spam catch' (or SC) rate refers to the catch rate on those corpora. The *Wombat* corpus did not contribute to the spam corpus, the calculation of the SC rate or the final score.

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 slightly less than 0.01%, while one missed email in the newsletter corpus results in an FP rate of about 0.5%).

The size of the *Wombat* corpus is also significantly smaller than the two spam corpora, and here too a single missed message results in a decrease in catch rate of almost 0.2%.

Bitdefender Security for Mail Servers 3.1.2

SC rate: 99.85%
FP rate: 0.03%
Final score: 99.68
Project Honey Pot SC rate: 99.86%
Abusix SC rate: 99.79%
Wombat SC rate: 58.8%
Newsletters FP rate: 0.0%

At VB2012, I attended a talk by one of *Bitdefender*'s anti-spam researchers on social networking spam. It was interesting, albeit slightly worrying, to see what possibilities there are for spammers. It was also good to see a VBSpam regular (*Bitdefender* remains the only participant to have submitted a product for every test) looking at other types of spam.

For this test, *Bitdefender* used a new machine, with *CentOS* rather than *SuSE* as the *Linux* flavour of choice, but the engine remained the same. Compared to the previous test, *Bitdefender* scraped off four-fifths of its false negative rate, making the four false positives (compared to zero in the previous test) almost understandable. The product earns its 22nd VBSpam award. The only concern was the rather low catch rate on the *Wombat* corpus, where the product blocked fewer than 60% of the emails; hopefully the next test will show that this was a one-off glitch.

ESET Mail Security for Microsoft Exchange Server

SC rate: 99.74%
FP rate: 0.00%
Final score: 99.74
Project Honey Pot SC rate: 99.71%
Abusix SC rate: 99.85%
Wombat SC rate: 90.1%
Newsletters FP rate: 9.4%

ESET takes great pride in its impressive record in our VB100 anti-malware tests and I don't expect

	True negatives	False positives	FP rate	False negatives	True positives	SC rate	Final score
Bitdefender	11996	4	0.03%	139	92027	99.85%	99.68
ESET	12000	0	0.00%	242	91924	99.74%	99.74
FortiMail	11999	1	0.01%	311	91855	99.66%	99.62
GFI	11998	2	0.02%	245	91921	99.73%	99.65
Halon Security	12000	0	0.00%	450	91716	99.51%	99.51
IBM	11996	4	0.03%	1232	90934	98.66%	98.49
Kaspersky LMS	11999	1	0.01%	140	92026	99.85%	99.81
Libra Esva	11999	1	0.01%	44	92122	99.95%	99.91
McAfee Email Gateway	11992	8	0.07%	240	91926	99.74%	99.41
McAfee SaaS	11997	3	0.03%	429	91737	99.53%	99.41
Netmail Secure	12000	0	0.00%	203	91963	99.78%	99.78
OnlyMyEmail	11999	1	0.01%	1	92165	99.999%	99.96
Sophos	11995	5	0.04%	819	91347	99.11%	98.90
SPAMfighter	11978	22	0.18%	553	91613	99.40%	98.48
SpamTitan	12000	0	0.00%	179	91987	99.81%	99.81
Symantec	11996	4	0.03%	338	91828	99.63%	99.46
The Email Laundry	11989	11	0.09%	359	91807	99.61%	99.15
Vamsoft ORF	11991	9	0.08%	624	91542	99.32%	98.95
ZEROSPAM	11997	3	0.03%	61	92105	99.93%	99.81
Spamhaus ZEN+DBL	12000	0	0.00%	5646	86520	93.87%	93.87
SURBL	12000	0	0.00%	35181	56985	61.83%	61.83

* *Spamhaus* and *SURBL* are both partial solutions and their performance is not to be compared with that of other products, neither should the performance of each be compared with the other.

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

the VBSpam tests to be any different. The anti-spam developers will thus be delighted with this month's test results, which see the product win a VBSpam+ award on only its second entry in the tests – the product had no false positives, and missed barely more than one in 400 spam emails.

Even the more difficult-to-filter spam was dealt with well, with the product blocking more than 90% of the *Wombat* corpus. Perhaps the only area where there is room for improvement is in the newsletter corpus – though of course, this doesn't contain emails that would be considered business critical.

Fortinet FortiMail

SC rate: 99.66%

FP rate: 0.01%

Final score: 99.62

Project Honey Pot SC rate: 99.58%

Abusix SC rate: 99.99%

Wombat SC rate: 79.0%

Newsletters FP rate: 4.0%

Only a single false positive stood in the way of *Fortinet* achieving a VBSpam+ award for its *FortiMail* appliance this month. Instead, its developers will have to settle for its 21st VBSpam award in as many tests.

The product's performance on the *Wombat* feed was slightly poorer than that of the average product thanks to a number of 'traditional' phishing emails being missed by the solution. Nevertheless, this is another very decent performance from what has become a VBSpam regular – and good motivation for its developers to work on reducing the FP rate that little bit further in the next test.

GFI MailEssentials

SC rate: 99.73%
FP rate: 0.02%
Final score: 99.65
Project Honey Pot SC rate: 99.73%
Abusix SC rate: 99.76%
Wombat SC rate: 88.6%
Newsletters FP rate: 10.4%

In the last test, *GFI*'s *MailEssentials* blocked more spam than in any previous test, and this time the product improved its catch rate once again, missing fewer than one in 375 spam messages.

At the same time, the product missed just two legitimate emails – one written in Dutch, the other in Finnish. Performance on the emails in the *Wombat* feed was average, and the fact that the product missed more newsletters than any other participating product is a minor concern – it would be nice if this could be fixed in the next test. Of course, this does not affect the final score, which earns the product its tenth VBSpam award.

Halon Security

SC rate: 99.51%
FP rate: 0.00%
Final score: 99.51
Project Honey Pot SC rate: 99.49%
Abusix SC rate: 99.58%
Wombat SC rate: 92.1%
Newsletters FP rate: 0.0%

I have regularly stressed the importance of products avoiding false positives altogether – which, especially given the increased size of the ham corpus this month, is no easy achievement. *Halon* managed just that however, as well as missing fewer than one in 200 spam emails. This impressive performance earns the Swedish solution its first VBSpam+ award.

Halon also deserves credit for being the only solution to avoid false positives both in the ham corpus and in the newsletter corpus. Finally, a very decent performance on the *Wombat* feed (one of the better scores in the test) is the icing on *Halon*'s cake.

IBM Lotus Protector for Mail Security

SC rate: 98.66%
FP rate: 0.03%
Final score: 98.49
Project Honey Pot SC rate: 98.45%
Abusix SC rate: 99.54%
Wombat SC rate: 78.6%
Newsletters FP rate: 0.0%

IBM suffered a bit of a glitch in the last test, with a rather high false positive rate. Thankfully, that proved to be a one-off, as the product missed only four legitimate emails in this month's enlarged ham corpus (and no newsletters).

Against that stands a bit of a drop in the product's spam catch rate: against the general tide, *IBM* missed more than one in 100 spam emails, and also seemed to struggle more than most with the *Wombat* feed of phishing emails. Still, the product did enough to achieve its eighth consecutive VBSpam award, and its developers will hopefully set to work on increasing the spam catch rate for next time.

Kaspersky Linux Mail Security 8.0

SC rate: 99.85%
FP rate: 0.01%
Final score: 99.81
Project Honey Pot SC rate: 99.82%
Abusix SC rate: 99.96%
Wombat SC rate: 74.0%
Newsletters FP rate: 0.0%

Two members of *Kaspersky*'s anti-spam team delivered presentations at VB2012, one of which gave some insight into how the product deals with new spam campaigns. I thought this was fascinating, and it was good to see that this method worked – the *Linux Mail Security* product had one of the highest spam catch rates in this test.

The false positive rate was better than average too, but with one missed legitimate email the developers will no doubt be a little disappointed to miss out on a VBSpam+ award by a whisker. Something that requires more serious improvement is the product's performance on the phishing emails in the *Wombat* corpus, where it missed more than one in four emails, but this shouldn't stop the developers from celebrating the product's third VBSpam award in as many tests – as well as the fourth highest final score.

Libra Esva 2.8

SC rate: 99.95%
FP rate: 0.01%
Final score: 99.91
Project Honey Pot SC rate: 99.94%
Abusix SC rate: 99.99%

	Newsletters		Project Honey Pot		Abusix		Wombat		pre-DATA†		STDev‡
	False positives	FP rate	False negatives	SC rate	False negatives	SC rate	False negatives	SC rate	False negatives	SC rate	
Bitdefender	0	0.0%	100	99.86%	39	99.79%	249	58.8%			0.37
ESET	19	9.4%	215	99.71%	27	99.85%	60	90.1%			0.38
FortiMail	8	4.0%	309	99.58%	2	99.99%	127	79.0%			0.47
GFI	21	10.4%	201	99.73%	44	99.76%	69	88.6%			0.41
Halon Security	0	0.0%	374	99.49%	76	99.58%	48	92.1%			0.69
IBM	0	0.0%	1148	98.45%	84	99.54%	129	78.6%			1.33
Kaspersky LMS	0	0.0%	133	99.82%	7	99.96%	157	74.0%			0.32
Libra Esva	1	0.5%	42	99.94%	2	99.99%	24	96.0%	84704	91.90%	0.16
McAfee Email Gateway	1	0.5%	169	99.77%	71	99.61%	65	89.2%			1.77
McAfee SaaS	6	3.0%	108	99.85%	321	98.23%	55	90.9%			1.84
Netmail Secure	20	9.9%	191	99.74%	12	99.93%	61	89.9%	81589	88.52%	0.35
OnlyMyEmail	4	2.0%	1	99.999%	0	100.00%	0	100.0%			0.03
Sophos	0	0.0%	403	99.46%	416	97.71%	38	93.7%			1.80
SPAMfighter	6	3.0%	532	99.28%	21	99.88%	0	100.0%			1.01
SpamTitan	1	0.5%	173	99.77%	6	99.97%	49	91.9%			0.38
Symantec	5	2.5%	268	99.64%	70	99.61%	164	72.9%			0.53
The Email Laundry	0	0.0%	321	99.57%	38	99.79%	35	94.2%	88040	95.52%	0.48
Vamsoft ORF	3	1.5%	591	99.20%	33	99.82%	181	70.0%			0.72
ZEROSPAM	15	7.4%	52	99.93%	9	99.95%	62	89.7%	87330	94.75%	0.19
Spamhaus ZEN+DBL*	0	0.0%	1921	97.40%	3725	79.47%	320	47.0%	84198	91.35%	3.99
SURBL*	0	0.0%	27504	62.84%	7677	57.69%	604	0.0%			13.87

* *Spamhaus* and *SURBL* are both partial solutions and their performance is not to be compared with that of other products, neither should the performance of each be compared with the other.

† pre-DATA filtering was optional and was applied on the full corpus. Four of the false positives for *The Email Laundry* and the single FP for *Libra Esva* occurred pre-DATA. The others were all 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.*)

Libra Esva 2.8 contd.

SC rate pre-DATA: 91.90%

Wombat SC rate: 96.0%

Newsletters FP rate: 0.5%

The developers of *Libra Esva* told me they were keen to earn another VBSpam+ award. They came very close to doing so, with the product's false positive rate dropping to just a single missed legitimate email. What is more impressive is that this was achieved while barely compromising on the spam catch rate, which at 99.95%, was the second highest in the test.

Difficult-to-filter spam is not a big problem for the virtual solution either – it blocked 96% of emails in the *Wombat* corpus, while in the newsletter feed there was only one false positive. With the second highest final score, even without the VBSpam+ award, there is plenty of reason for *Libra* to celebrate its 16th consecutive VBSpam award.

McAfee Email Gateway 7.0

SC rate: 99.74%

FP rate: 0.07%

Final score: 99.41

Project Honey Pot SC rate: 99.77%

McAfee Email Gateway 7.0 contd.

Abusix SC rate: 99.61%
Wombat SC rate: 89.2%
Newsletters FP rate: 0.5%

After a continuous increase throughout the year, *McAfee*'s *Email Gateway* appliance saw its spam catch rate drop a little in this test – but if that's the price to be paid for a decreased false positive rate, it's a price worth paying, especially since the product still only missed a little more than one in 400 spam messages.

The product's performance on both the *Wombat* feed and the newsletters (where it missed just a single message) was good, and another VBSpam award was easily achieved – hopefully the developers will be motivated to see if the FP rate can be reduced a little further.

McAfee SaaS Email Protection

SC rate: 99.53%
FP rate: 0.03%
Final score: 99.41
Project Honey Pot SC rate: 99.85%
Abusix SC rate: 98.23%
Wombat SC rate: 90.9%
Newsletters FP rate: 3.0%

Like the company's appliance, *McAfee*'s hosted *SaaS Email Protection* product saw a decrease in both its spam catch rate and its false positive rate. Assuming the drop in the former isn't too big, I tend to see this as a good thing and indeed, the product wins its ninth VBSpam award with a slightly increased final score.

There were a handful of false positives in the newsletter corpus, but it was nice to see that the hosted solution would have prevented more than 90% of the phishing emails in the *Wombat* corpus from making it to users' inboxes.

Messaging Architects Netmail Secure

SC rate: 99.78%
FP rate: 0.00%
Final score: 99.78
Project Honey Pot SC rate: 99.74%
Abusix SC rate: 99.93%
SC rate pre-DATA: 88.52%
Wombat SC rate: 89.9%
Newsletters FP rate: 9.9%

As a rule of thumb, there is a direct correlation between a product's spam catch rate and its false positive rate: an increase or decrease in one has the same effect on the other. So it was nice to see *Netmail* make an exception to this rule: while the product's catch rate increased, its false positive rate dropped. What is more, it dropped to zero, thus earning the product its first VBSpam+ award.

The appliance from *Messaging Architects* blocked almost 90% of the emails in the *Wombat* feed, with the 20 false positives in the newsletter corpus the only small concern.

OnlyMyEmail's Corporate MX-Defender

SC rate: 99.999%
FP rate: 0.01%
Final score: 99.96
Project Honey Pot SC rate: 99.999%
Abusix SC rate: 100.00%
Wombat SC rate: 100.00%
Newsletters FP rate: 2.0%

OnlyMyEmail's 99.999% catch rate on the spam corpus won't come as a surprise for regular *VB* readers: in this test, the only spam message the product missed was a spammy newsletter (one that was missed by many products and that, incidentally, made the amusing mistake of 'including' images by referencing their disk location). What should impress readers is that the product had no problems with the *Wombat* feed either – which, judging by other products' performance, is significantly more difficult to filter – none of the 604 emails were missed by the hosted solution.

The product missed one legitimate email, and was thus denied a VBSpam+ award, and also missed four newsletters. However, this shouldn't distract from the product winning its 13th consecutive VBSpam award, once again doing so with the highest final score.

Sophos Email Appliance

SC rate: 99.11%
FP rate: 0.04%
Final score: 98.90
Project Honey Pot SC rate: 99.46%
Abusix SC rate: 97.71%
Wombat SC rate: 93.7%
Newsletters FP rate: 0.0%

Sophos's *Email Appliance* saw a decrease in its spam catch rate – the product had some issues with predominantly Chinese spam in our corpus. However, I should note that a significant chunk of the emails the products missed were sent from a single IP address – if a real organization suffered such an attack, it is likely

Hosted solutions	Anti-malware	IPv6	DKIM	Multiple MX-records	Multiple locations
McAfee SaaS	McAfee	√	√	√	√
OnlyMyEmail	Proprietary (optional)		√	√	√
The Email Laundry	Included*		√	√	√
ZEROSPAM	ClamAV			√	√

* Vendor prefers not to reveal identity of anti-malware engine. (*Please refer to the text for full product names.*)

Local solutions	Anti-malware	IPv6	DKIM	Interface			
				CLI	GUI	Web GUI	API
Bitdefender	Bitdefender	√		√		√	
ESET	ESET Threatsense			√	√		
FortiMail	Fortinet	√	√	√		√	
GFI	Five anti-virus engines	√				√	
Halon Security	Commtouch, Kaspersky, ClamAV, HRPS	√	√	√		√	√
IBM	Sophos			√		√	
Kaspersky LMS	Kaspersky	√					
Libra Esva	ClamAV; others optional		√			√	
McAfee Email Gateway	McAfee	√	√	√	√	√	
Netmail	Proprietary	√	√	√			
Sophos	Sophos					√	
SPAMfighter	VIRUSfighter (optional)	√	√			√	
SpamTitan	Kaspersky, ClamAV	√	√	√		√	√
Symantec	Symantec		√	√		√	
Vamsoft ORF	Optional*				√		

* Various engines can be plugged in. (*Please refer to the text for full product names.*)

that a system administrator would block the IP, which is something that many products (including *Sophos*) allow.

There was better news regarding the product's performance on the *Wombat* corpus, where *Sophos* missed only one in 16 emails, outperforming many of its competitors. There were five false positives (but none on the newsletters) and the product earns it 17th VBSpam award.

SPAMfighter Mail Gateway

SC rate: 99.40%
FP rate: 0.18%
Final score: 98.48
Project Honey Pot SC rate: 99.28%
Abusix SC rate: 99.88%
Wombat SC rate: 100.0%
Newsletters FP rate: 3.0%

This month's test results are a bit of a mixed bag for *SPAMfighter*. On the one hand, it did prove that last month's drop in spam catch rate was a one-off glitch, as the product blocked well over 99% of spam this time (including a message that claimed to come from spamfighter.com – perhaps a sign that the product's free home-user version has become rather well known). More impressively, the product didn't miss any of the 604 emails in the *Wombat* corpus, which is no trivial achievement given that several products blocked fewer than 80% of these emails.

However, the product missed 22 legitimate emails – more than any other product. This did not prevent it from winning its 12th VBSpam award, as the final score still exceeded the benchmark 97. Nevertheless, it would be nice to have some reassurance in the next test that a high FP rate is not the price that has to be paid for the rebound in SC rate.

SpamTitan 5.11

SC rate: 99.81%
FP rate: 0.00%
Final score: 99.81
Project Honey Pot SC rate: 99.77%
Abusix SC rate: 99.97%
Wombat SC rate: 91.9%
Newsletters FP rate: 0.5%

This test saw *SpamTitan* miss fewer than one in 500 spam emails, while at the same time making the correct decision on all 12,000 legitimate emails. This impressive performance earns the product its second VBSpam+ award – and *SpamTitan* becomes the first product to receive more than one VBSpam+ award.

The product managed a decent catch rate on the *Wombat* corpus and only missed one newsletter, as well as achieving this month's third-highest final score.

Symantec Messaging Gateway 10.0

SC rate: 99.63%
FP rate: 0.03%
Final score: 99.46
Project Honey Pot SC rate: 99.64%
Abusix SC rate: 99.61%
Wombat SC rate: 72.9%
Newsletters FP rate: 2.5%

The spam catch rate for *Symantec*'s virtual appliance saw a small drop in this test – though at 99.63% this is hardly a serious problem. Thankfully, the product's false positive rate remained low at 0.03% (and all false positives were from the same sender) and thus the security giant wins its 18th consecutive VBSpam award.

Messaging Gateway missed five newsletters, while performance on the *Wombat* feed was a little below par – the product missed a fairly large number of typical phishing emails targeting a wide range of banks. It would be nice to see this improved upon in future tests.

The Email Laundry

SC rate: 99.61%
FP rate: 0.09%
Final score: 99.15
Project Honey Pot SC rate: 99.57%
Abusix SC rate: 99.79%
SC rate pre-DATA: 95.52%
Wombat SC rate: 94.2%
Newsletters FP rate: 0.0%

Complete solutions sorted by final score	
OnlyMyEmail's Corporate MX-Defender	99.96
Libra Esva 2.8	99.91
SpamTitan 5.11	99.81
Kaspersky Linux Mail Security 8.0	99.81
ZEROSPAM	99.81
Messaging Architects Netmail Secure	99.78
ESET Mail Security for MS Exchange Server	99.74
Bitdefender Security for Mail Servers 3.1.2	99.68
GFI MailEssentials	99.65
Fortinet FortiMail	99.62
Halon Security	99.51
Symantec Messaging Gateway 10.0	99.46
McAfee Email Gateway 7.0	99.41
McAfee SaaS Email Protection	99.41
The Email Laundry	99.15
Vamsoft ORF	98.95
Sophos Email Appliance	98.90
IBM Lotus Protector for Mail Security	98.49
SPAMfighter Mail Gateway	98.48

The Email Laundry saw a slight deterioration in its performance in this test, with both a small decrease in its spam catch rate and a small increase in its false positive rate (which, at 0.09%, was the second highest this month). Although pre-DATA catch rates have fallen this year for all products where we measure them, *The Email Laundry* did manage to outperform its competitors here. The product's final score was well over 99 and thus the hosted solution easily earns another VBSpam award despite the drop in performance.

Performance on the other feeds was good, with no false positives in the newsletter corpus, while the product caught more than 94% of the spam in the *Wombat* corpus.

Vamsoft ORF

SC rate: 99.32%
FP rate: 0.08%
Final score: 98.95
Project Honey Pot SC rate: 99.20%
Abusix SC rate: 99.82%
Wombat SC rate: 70.0%
Newsletters FP rate: 1.5%

The re-introduction of the *Abusix* feed in the last test caused some problems for *Vamsoft ORF*, which

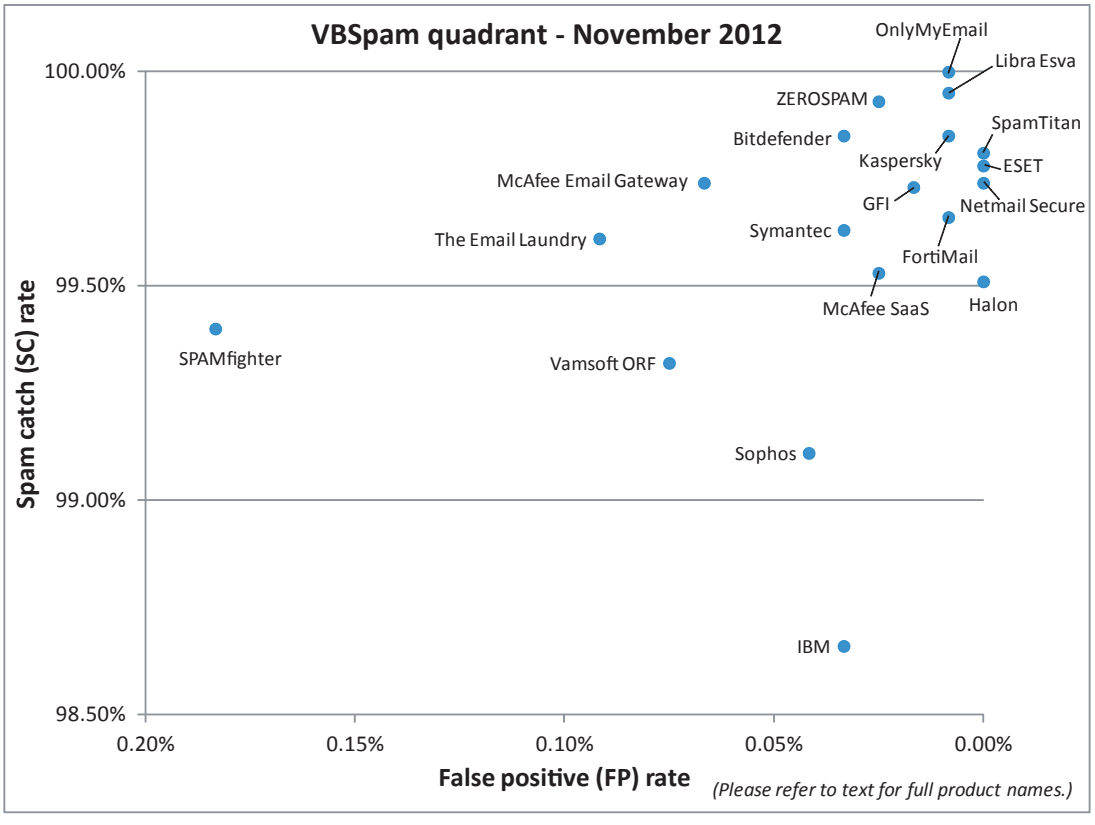

narrowly missed out on a VBSpam award. It was thus good to see the product's performance bounce back, blocking well over 99% of all spam. Against that stood an increase in the false positive rate – traditionally the product's strong point – with the product missing nine emails from a wide range of senders. This didn't jeopardize its chances of a VBSpam award though, as the final score was well within the margins to earn one once again.

Three newsletters from an Irish newspaper were missed by *ORF*, while its performance on the *Wombat* feed stood at 70% – it would be nice to see an improvement in this in the next test.

ZEROSPAM

SC rate: 99.93%
FP rate: 0.03%
Final score: 99.81
Project Honey Pot SC rate: 99.93%
Abusix SC rate: 99.95%
SC rate pre-DATA: 94.75%
Wombat SC rate: 89.7%
Newsletters FP rate: 7.4%

Since joining the VBSpam tests at the beginning of the year *ZEROSPAM* has come very close to living up to its name, generally missing very few spam messages. This test was no exception as the hosted solution blocked more spam than all but two other products. There were three false positives, and the product easily earns its fifth VBSpam award.

The Canadian product blocked almost 90% of the phishing messages in the *Wombat* corpus, while there were 15 false positives on newsletters in a number of different languages – if anything, it would be nice to see this rate reduced a little.

Spamhaus ZEN+DBL

SC rate: 93.87%
FP rate: 0.00%
Final score: 93.87
Project Honey Pot SC rate: 97.40%
Abusix SC rate: 79.47%
SC rate pre-DATA: 91.35%
Wombat SC rate: 47.0%
Newsletters FP rate: 0.0%

The inclusion of *Spamhaus* in these test results gives a good picture of the spam landscape, as it shows as much about

the quality of the IP and domain blacklists as it does about the spammers' tendency to use compromised legitimate (and thus generally unblockable) domains and IP addresses.

Given the performance of those products for which we measure pre-DATA catch rates, it is fair to say that it's mostly the spammers' increased use of legitimate servers that has caused the drop in *Spamhaus*'s catch rate. Still, more than 15 out of 16 emails were blocked by either the *ZEN* aggregated IP-blacklist or the *DBL* domain-blacklist, without any further inspection of body and headers.

It was nice to see the product block almost half of the emails in the *Wombat* feed, while there were no false positives in either the ham corpus or the newsletter corpus.

SURBL

SC rate: 61.83%
FP rate: 0.00%
Final score: 61.83
Project Honey Pot SC rate: 62.84%
Abusix SC rate: 57.69%
Wombat SC rate: 0.0%
Newsletters FP rate: 0.0%

The fact that *SURBL* did not block any messages in the *Wombat* corpus probably says a lot about the corpus and confirms my belief that the creators of these emails take care not to use blacklisted domains. The *SURBL* domain blacklist blocked almost 62% of the emails in the spam corpus based solely on the presence of a blacklisted domain, while as usual it did not miss a single legitimate email.

CONCLUSION

The significant increase in products' spam catch rates is undeniably good news, especially given the decline we saw earlier this year. Extrapolating on the four-cents-per-spam cost I mentioned in the introduction, this means that in the past couple of months, spam filters will have saved organizations thousands of dollars more than they would have done with this spring's performance.

However, the increase in some of the products' false positive rates shows there is still work to be done, while the performance on the *Wombat* feed shows that a product that blocks more than 99 out of 100 spam messages doesn't necessarily do the same when it comes to trickier and potentially more harmful spam.

The next VBSpam test will run in December 2012, with the results scheduled for publication in January. Developers interested in submitting products should email martijn.grooten@virusbtn.com.

VIRUS BULLETIN

Editor: Helen Martin

Technical Editor: Dr Morton Swimmer

Test Team Director: John Hawes

Anti-Spam Test Director: Martijn Grooten

Security Test Engineer: Simon Bates

Sales Executive: Allison Sketchley

Perl Developer: Tom Gracey

Consulting Editors:
Nick FitzGerald, *AVG, NZ*
Ian Whalley, *Google, USA*
Dr Richard Ford, *Florida Institute of Technology, USA*

SUBSCRIPTION RATES

Subscription price for Virus Bulletin magazine (including comparative reviews) for 1 year (12 issues):

- Single user: $175
- Corporate (turnover < $10 million): $500
- Corporate (turnover < $100 million): $1,000
- Corporate (turnover > $100 million): $2,000
- *Bona fide* charities and educational institutions: $175
- Public libraries and government organizations: $500

Corporate rates include a licence for intranet publication.

Subscription price for Virus Bulletin comparative reviews only for 1 year (6 VBSpam and 6 VB100 reviews):

- Comparative subscription: $100

See http://www.virusbtn.com/virusbulletin/subscriptions/ for subscription terms and conditions.

Editorial enquiries, subscription enquiries, orders and payments:

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/