COMPARATIVE REVIEW

VBSPAM COMPARATIVE – MAY 2011

Martijn Grooten

When we embarked on VBSpam testing two years ago we had a number of goals, one of which was to provide the anti-spam community – both the developers and the users of spam filters – with information that would help them in the fight against unwanted email.

I still find this one of the most rewarding aspects of my job: developers let me know that they have used feedback from our tests to tweak their filters in order to provide better protection for their customers' inboxes. But with reward comes a responsibility: we have to make sure that the feedback we give is relevant and reflects a real situation.

I have been asked several times why we don't add a stream of newsletters to the test. Although we would like to do this (and, in fact, hope to do so soon), we are very hesitant about incorporating newsletter filtering performance into the criteria for earning a VBSpam award.

There are few email users who do not subscribe to at least a small number of newsletters. But there are even fewer users who have not received unwanted and apparently unsolicited newsletters. These messages may be straightforward spam. However, as a result of legal loopholes and small print in terms and conditions, they may not be spam according to some definitions. In some instances the user may even have subscribed to the newsletters and then forgotten they had done so.

This does not change the user's experience of receiving unwanted email or spam – but there is a very real possibility that exactly the same newsletters are *wanted* by other users. For this reason, we cannot make absolute statements about the right or wrong way to treat newsletters.

As a consequence, the performance numbers we currently report – in particular the spam catch rates – could well be better than those experienced by an organization using the product in a real-world situation. We do not think this should matter. After all, mail traffic can differ greatly between organizations, and so will spam catch rates and false positive rates. Moreover, to fully understand the meaning of a catch rate, one has to know the size of the full inbound mail stream; few people do.

What does matter, though, is the fact that our tests are comparative. This doesn't just mean that we test multiple products under the same circumstances, it also means that the results can and should be compared. A product that blocks 99.70% of spam in our tests might not have the same catch rate when used by a customer, but it is likely to perform better for that customer than a product that only catches 99.10% of spam in our tests.

The 13th VBSpam test did not prove unlucky for any of the participating products, with all 19 entrants receiving a VBSpam award. This report presents the detailed figures that distinguish the good products from the really good ones, but perhaps the most important question for potential customers to ask is: why did other products opt not to be tested?

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. Three 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 *SuSE Linux Enterprise Server 11*; the *Windows Server* products ran on either the 2003 or the 2008 version, depending on which was recommended by the vendor.

Two of the virtual products tested ran on *VMware ESXi 4.1*, while two others ran on *VMware Server 2.0*, which we had hitherto used for all virtual products.

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:

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

THE EMAIL CORPUS

The test ran for 16 consecutive days, from 12am GMT on Saturday 9 April 2011 until 12am GMT on Monday 25 April 2011.

The corpus contained 74,746 emails, 72,008 of which were spam. Of these, 40,674 were provided by *Project Honey Pot* and 31,334 were provided by *Abusix*; in both cases, the messages were relayed in real time, as were the 2,738 legitimate emails. As before, the legitimate emails were sent in a number of languages to represent an international mail stream and came from countries all over the world, including India, Russia, El Salvador and Japan.

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

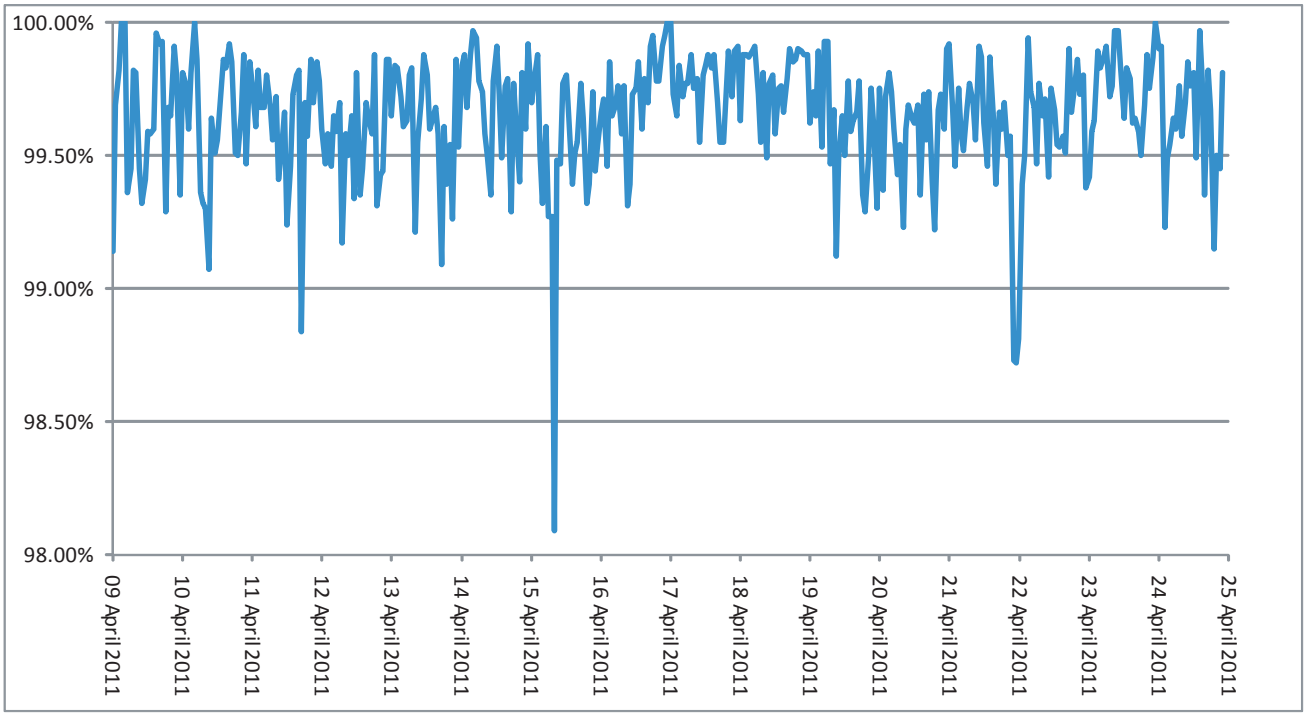

Figure 1: Average spam catch rate of all full solutions throughout the test.

As before, we looked at differences between the full corpus and the subset of 'difficult' spam – defined as those messages missed by at least two different filters; the latter concerned slightly more than 1 in 36 messages, which is a higher ratio than during previous tests.

This time we looked at the uniqueness of the messages in the full corpus of spam. An important characteristic of the vast majority of spam is that it is sent in bulk. Many filters make use of this characteristic to detect spam and it is commonly believed that making spam more 'unique' increases the likelihood that it will stay under the radar and remain uncaught.

We used a rather simple definition of uniqueness, determining two messages to be 'similar' if their subjects were equal[1] and if the number of lines in the raw bodies were equal. Of course, this is by no means the best way to define uniqueness, and an important part of anti-spam research is about finding ways to group similar messages. Still, it is an easy way to test the commonly held belief that unique messages are harder to filter.

In Table 1, on the left it can be seen that 22.2% of all spam was unique according to the above definition. 12.9% of all spam was part of a group of two to five similar messages, and so on; finally, 52.4% of all spam was part of a group of 51 or more messages (the largest group consisted of 615 pill spam messages). On the right-hand side, it is shown that among the 'difficult' spam, 28.5% was unique (within the full corpus), and 38.1% of difficult spam was part of a group of 51 or more messages – indicating that there is some truth in the belief that unique messages are harder to filter.

No. of similar messages	Percentage of spam corpus
1	22.2%
2 to 5	12.9%
6 to 10	3.1%
11 to 20	2.6%
21 to 50	6.8%
51+	52.4%

No. of similar messages	Percentage of 'difficult' spam
1	28.5%
2 to 5	10.5%
6 to 10	4.7%
11 to 20	5.4%
21 to 50	12.8%
51+	38.1%

Table 1: Left: Uniqueness of messages seen in the spam feeds. Right: Uniqueness of spam messages missed by at least two full solutions.

By using more advanced definitions of similarity and a larger corpus, one would be able to obtain more refined

[1] We only used one recipient in the test; hence messages with the recipient's local-part, domain or email address in the subject could still be considered similar.

results and possibly show a stronger correlation between uniqueness and difficulty of filtering.

BLACK- AND WHITELISTING

In our tests, we have always focused on the core of the spam filter: the anti-spam engine.

However, a spam filter comprises much more than its engine. This fact was brought home recently to security company *RSA*, one of whose employees received a targeted spam message containing a malicious attachment. The company's spam filter did its job, putting the message into quarantine; however, this did not stop the user from fishing the email out of quarantine, opening the attachment and thus exposing a backdoor into the company's network.

In this and future tests, we will look at a number of different features of spam filters, reporting on whether the products on test have these features and, if they do, whether they work as intended. We start this month by looking at black- and whitelisting.

We considered four possible features:

- The possibility to whitelist email coming from a certain IP address.
- The possibility to blacklist email coming from a certain IP address.
- The possibility to whitelist email based on the senders' domain[2].
- The possibility to blacklist email based on the senders' domain.

It is very important to note that we do not wish to make assertions as to whether this is something spam filters *should* have. Several participating filters do not provide some or all of these features, and they may have good reasons for not doing so. Using a black- or whitelist incorrectly could lead to many missed emails, or to inboxes overflowing with spam.

Still, for many a system administrator having these options may be just what is needed to allow mail through from an organization whose messages keep being blocked, or to stop newsletters that fail to respond to unsubscribe requests.

We tested these properties by slightly modifying emails that were previously blocked (to test whitelisting) or allowed (to test blacklisting) and resending them. The modifications to the products' settings were made by the testers following the developers' instructions.

[2] We tested by sending emails where both the MAIL FROM address in the SMTP envelope and the From: address in the email headers were on the white- or blacklisted domain.

RESULTS

AnubisNetworks Mail Protection Service

SC rate: 99.92%

FP rate: 0.07%

Final score: 99.55

Project Honey Pot SC rate: 99.88%

Abusix SC rate: 99.96%

AnubisNetworks' R&D recently built a tool[3] that tracks *Twitter* spammers, to demonstrate the company's awareness of the fact that unwanted messages are not restricted to email. This didn't detract from the performance of the spam filter though – the product blocked close to 100 per cent of spam messages and two false positives (the first since September last year) were not enough to get in the way of winning a sixth VBSpam award.

Both black- and whitelisting are possible with this product. The options were easily set in the web interface and worked very well.

BitDefender Security for Mail Servers 3.0.2

SC rate: 99.84%

FP rate: 0.00%

Final score: 99.84

Project Honey Pot SC rate: 99.85%

Abusix SC rate: 99.83%

BitDefender's anti-spam product runs on a number of platforms, but we have been testing the *Linux* product – integrated with *Postfix* – since the very first test. *Linux* fans will be pleased to learn that domain black- and whitelists can be activated by adding the domains to a configuration file. IP black- and whitelisting is not possible, but as the product is an SMTP proxy, blacklisting is usually an option in the ambient SMTP server.

What should please users of the product even more is that it earned yet another VBSpam award and continues to be the only product to have won an award in all 13 VBSpam tests. Moreover, with an excellent spam catch rate and zero false positives, its final score was in the top five of this test.

[3] http://www.tweetspike.org/.

eleven eXpurgate Managed Service 3.2

SC rate: 99.64%
FP rate: 0.00%
Final score: 99.64
Project Honey Pot SC rate: 99.43%
Abusix SC rate: 99.92%

eleven saw a slight improvement in its spam catch rate, but probably more important to the developers is the fact that it did not generate any false positives in this test – and has not done so in its last three tests. The company can thus pride itself on a third VBSpam award.

The web interface that comes with the hosted solution allows for black- and whitelisting of both IPs and domains, and setting them was a trivial task.

Fortinet FortiMail

SC rate: 99.82%
FP rate: 0.00%
Final score: 99.82
Project Honey Pot SC rate: 99.76%
Abusix SC rate: 99.91%

A regular VBSpam participant and repeated award winner, this month *Fortinet* wins its 12th consecutive VBSpam award. *FortiMail* equalled last month's spam catch rate and once again achieved a zero false positive rate, giving it the same high final score.

The web interface that is used to control the appliance allows for both black- and whitelisting; setting it up and getting it to work presented few problems.

GFI MailEssentials

SC rate: 99.62%
FP rate: 0.51%
Final score: 97.07
Project Honey Pot SC rate: 99.52%
Abusix SC rate: 99.76%

GFI is not new to the VBSpam tests, having won several VBSpam awards with *VIPRE*, a product originally developed by *Sunbelt*. However, *MailEssentials* is the Maltese company's own in-house-developed product. It runs on *Windows* and hooks into a number of SMTP servers, including *Exchange* and *IIS*. We tested it using the latter.

Set-up was easy and the user interface is clear and intuitive. It can be used to manage domain black- and whitelists and IP whitelists, all of which worked in a pretty straightforward manner. The product does not allow for IP blacklisting, but the developers recommend customers use the ambient SMTP server for that.

The product did well at catching spam, but generated twice as many false positives as the next most poorly performing product. *MailEssentials* did just scrape a high enough final score to win a VBSpam award, but it will now be up to *GFI*'s developers to show that the high false positive rate can be lessened by making some modifications to the product and/or its settings.

Halon Mail Security

SC rate: 99.46%
FP rate: 0.00%
Final score: 99.46
Project Honey Pot SC rate: 99.78%
Abusix SC rate: 99.06%

I have said it might take a test or two for products to fully adapt to our set-up and environment, but *Halon Mail Security* proved me wrong in the last test. The Swedish product (we tested the virtual appliance) combined a rather good spam catch rate with zero false positives on its first entry, and repeated the achievement in this test, winning its second VBSpam award.

Unsurprisingly, given the scripting language that can be used to tweak the product, the addition of black- and whitelists is possible – and these checks can take place at different places during the transaction. We chose the most straightforward places (less tech-savvy users will be pleased to know that no scripting was involved) and found them to work well.

Kaspersky Anti-Spam 3.0

SC rate: 99.37%
FP rate: 0.00%
Final score: 99.37
Project Honey Pot SC rate: 99.46%
Abusix SC rate: 99.25%

	True negatives	False positives	FP rate	False negatives	True positives	SC rate	Final score
AnubisNetworks	2736	2	0.07%	58	71950	99.92%	99.55
BitDefender	2738	0	0.00%	117	71891	99.84%	99.84
eleven	2738	0	0.00%	256	71752	99.64%	99.64
FortiMail	2738	0	0.00%	127	71881	99.82%	99.82
GFI MailEssentials	2724	14	0.51%	271	71737	99.62%	97.07
Halon Mail Security	2738	0	0.00%	386	71622	99.46%	99.46
Kaspersky Anti-Spam	2738	0	0.00%	454	71554	99.37%	99.37
Libra Esva	2738	0	0.00%	40	71968	99.94%	99.94
McAfee Email Gateway	2738	0	0.00%	57	71951	99.92%	99.92
McAfee EWS	2737	1	0.04%	1054	70954	98.54%	98.35
OnlyMyEmail	2738	0	0.00%	2	72006	100.00%	100.00
Sophos Email Appliance	2737	1	0.04%	169	71839	99.77%	99.58
SPAMfighter	2734	4	0.15%	213	71795	99.70%	98.97
SpamTitan	2731	7	0.26%	46	71962	99.94%	98.66
Symantec Messaging Gateway	2738	0	0.00%	74	71934	99.90%	99.90
The Email Laundry	2736	2	0.07%	168	71840	99.77%	99.40
Vade Retro	2736	2	0.07%	1336	70672	98.14%	97.78
Vamsoft ORF	2738	0	0.00%	417	71591	99.42%	99.42
Spamhaus*	2738	0	0.00%	799	71209	98.89%	98.89

*As the only partial solution in this test, the results for *Spamhaus* are listed separately from the full solutions.
(Please refer to text for full product names.)

Since installing *Kaspersky Anti-Spam* two years ago, I have not needed to look at the web interface and I had almost even forgotten what kind of interface the product used. That, of course, is a good thing as it demonstrates that the product has been running without issues. On revisiting the interface, to add black- and whitelists, I found it to be intuitive and easily navigated.

It was good to see that last month's drop in performance was a one-off affair, with the spam catch rate returning to well over 99%, still with no false positives. This performance easily wins *Kaspersky* its 11th VBSpam award.

Libra Esva 2.0

SC rate: 99.94%
FP rate: 0.00%
Final score: 99.94
Project Honey Pot SC rate: 99.95%
Abusix SC rate: 99.94%
SC rate pre-DATA: 98.64%

Prior to this test, we moved *Libra Esva*'s virtual product to our new *VMware ESXi 4.1* server. The move was easy, thanks to good work from both the product's and *VMware*'s developers. The product's simple web interface allows for the

use of IP and domain black- and whitelists, but domain whitelisting takes place after the SMTP traffic is checked against the IP blacklists used by the product, and thus might not work in all cases.

Users may have little need for domain whitelisting though. With zero false positives and a 99.94% spam catch rate, *Libra Esva* wins its seventh VBSpam with the second highest final score for the third time in a row.

McAfee Email Gateway (formerly IronMail)

SC rate: 99.92%
FP rate: 0.00%
Final score: 99.92
Project Honey Pot SC rate: 99.87%
Abusix SC rate: 99.98%

Like most products, *McAfee*'s *Email Gateway* appliance is controlled by a web interface. I had not looked at the interface for some time, and was impressed by its many bells and whistles. Black- and whitelisting are possible, though less straightforward than with most products. However, given the huge consequences mistakes can have, making sure users really know what they're doing might not be a bad thing.

Of course, if the product performs its main task well there should be little need for black- and whitelisting. This is certainly the case for the *Email Gateway* appliance, and with a spam catch rate of 99.92% and zero false positives (down from six), it wins its 11th consecutive VBSpam award with the third highest final score.

McAfee Email and Web Security Appliance

SC rate: 98.54%
FP rate: 0.04%
Final score: 98.35
Project Honey Pot SC rate: 98.73%
Abusix SC rate: 98.28%

I've always been charmed by the web interface of *McAfee*'s *EWS* appliance with its many options and, unsurprisingly, it let me add black- and whitelists easily.

However, the product's developers will be more concerned with its performance, given that it failed to win a VBSpam award in the last test. Happily, *EWS*'s performance improved significantly – this time generating only a single false positive. There is still some room for improvement, but the product nevertheless wins its tenth VBSpam award.

OnlyMyEmail's Corporate MX-Defender

SC rate: 100.00%
FP rate: 0.00%
Final score: 100.00
Project Honey Pot SC rate: 100.00%
Abusix SC rate: 99.99%

To use black- and whitelists, users of *OnlyMyEmail*'s *MX-Defender* need to fill out a form on the company's website. I received a quick response to this and soon realised the benefit of a second check of my request: I had made a mistake in filling out the form and was asked if I really wanted what I had asked for. After clarification, the black- and whitelists were added. The IP whitelist does not always work though, and blocked a very spammy test message from the whitelisted IP – which hints towards the fact that pure whitelisting is not always a good idea.

But with zero false positives, *OnlyMyEmail*'s users will find little need to use whitelisting. And they will not need to add many items (if any) to the blacklist either, as the product missed just two out of over 70,000 spam messages. In fact, not only did the product achieve the highest spam catch rate for the fourth time in a row, it also achieved the highest final score in the test for the second time and, rounded to 100, the highest final score and spam catch rate since our tests began two years ago. Needless to say, *OnlyMyEmail* can be extremely proud of its fourth VBSpam award.

Sophos Email Appliance

SC rate: 99.77%
FP rate: 0.04%
Final score: 99.58
Project Honey Pot SC rate: 99.77%
Abusix SC rate: 99.75%

In March, *Sophos* won its eighth VBSpam award in as many tests and this month it easily adds a ninth to its tally. A slight drop in performance is not a serious issue for the appliance – which achieved the highest final score in the previous test.

The simple web interface allows for both black- and whitelisting by domain and IP address and all worked as expected.

	Project Honeypot		Abusix		pre-DATA†						
	False negative	SC rate	False negative	SC rate	False negative	SC rate	STDev‡	IP WL	IP BL	Dom WL	Dom BL
AnubisNetworks	47	99.88%	11	99.96%			0.21	+	+	+	+
BitDefender	63	99.85%	54	99.83%			0.39	-	-	+	+
eleven	230	99.43%	26	99.92%			0.94	+	+	+	+
FortiMail	99	99.76%	28	99.91%			0.36	+	+	+	+
GFI MailEssentials	197	99.52%	74	99.76%			0.49	+	-	+	+
Halon Mail Security	91	99.78%	295	99.06%			0.76	+	+	+	+
Kaspersky Anti-Spam	218	99.46%	236	99.25%			1.12	+	+	+	+
Libra Esva	22	99.95%	18	99.94%	28949	98.64%	0.17	+	+	+	+
McAfee Email Gateway	51	99.87%	6	99.98%			0.21	+	+	+	+
McAfee EWS	515	98.73%	539	98.28%			1.48	+	+	+	+
OnlyMyEmail	0	100.00%	2	99.99%			0.04	-	+	+	+
Sophos Email Appliance	92	99.77%	77	99.75%			0.51	+	+	+	+
SPAMfighter	106	99.74%	107	99.66%			0.57	+	+	+	+
SpamTitan	20	99.95%	26	99.92%			0.19	-	+	+	+
Symantec Messaging Gateway	56	99.86%	18	99.94%			0.23	+	+	+	+
The Email Laundry	135	99.67%	33	99.89%	29249	99.30%	0.35	-	-	+	+
Vade Retro	352	99.13%	984	96.86%			2.51	-	-	-	-
Vamsoft ORF	320	99.21%	97	99.69%			0.60	+	+	+	+
Spamhaus*	557	98.63%		99.23%	30740	98.15%	1.05	-	-	-	-

† pre-DATA filtering was optional and was applied on the full spam corpus. All of *The Email Laundry*'s false positives occurred pre-DATA; none of the other products had pre-DATA false positives.

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

* As the only partial solution in this test, the results for *Spamhaus* are listed separately from the full solutions.

(Please refer to text for full product names.)

SPAMfighter Mail Gateway

SC rate: 99.70%

FP rate: 0.15%

Final score: 98.97

Project Honey Pot SC rate: 99.74%

Abusix SC rate: 99.66%

The simple web interface provided for system administrators to modify *SPAMfighter*'s settings made it a simple process to find and add black- and whitelists.

Not only do the developers deserve praise for that, but even more so for the fact that the product achieved its highest spam catch rate to date. With just four false positives, this also gave the product its highest final score and *SPAMfighter* earns its tenth VBSpam award.

SpamTitan

SC rate: 99.94%

FP rate: 0.26%

Final score: 98.66

Project Honey Pot SC rate: 99.95%

Abusix SC rate: 99.92%

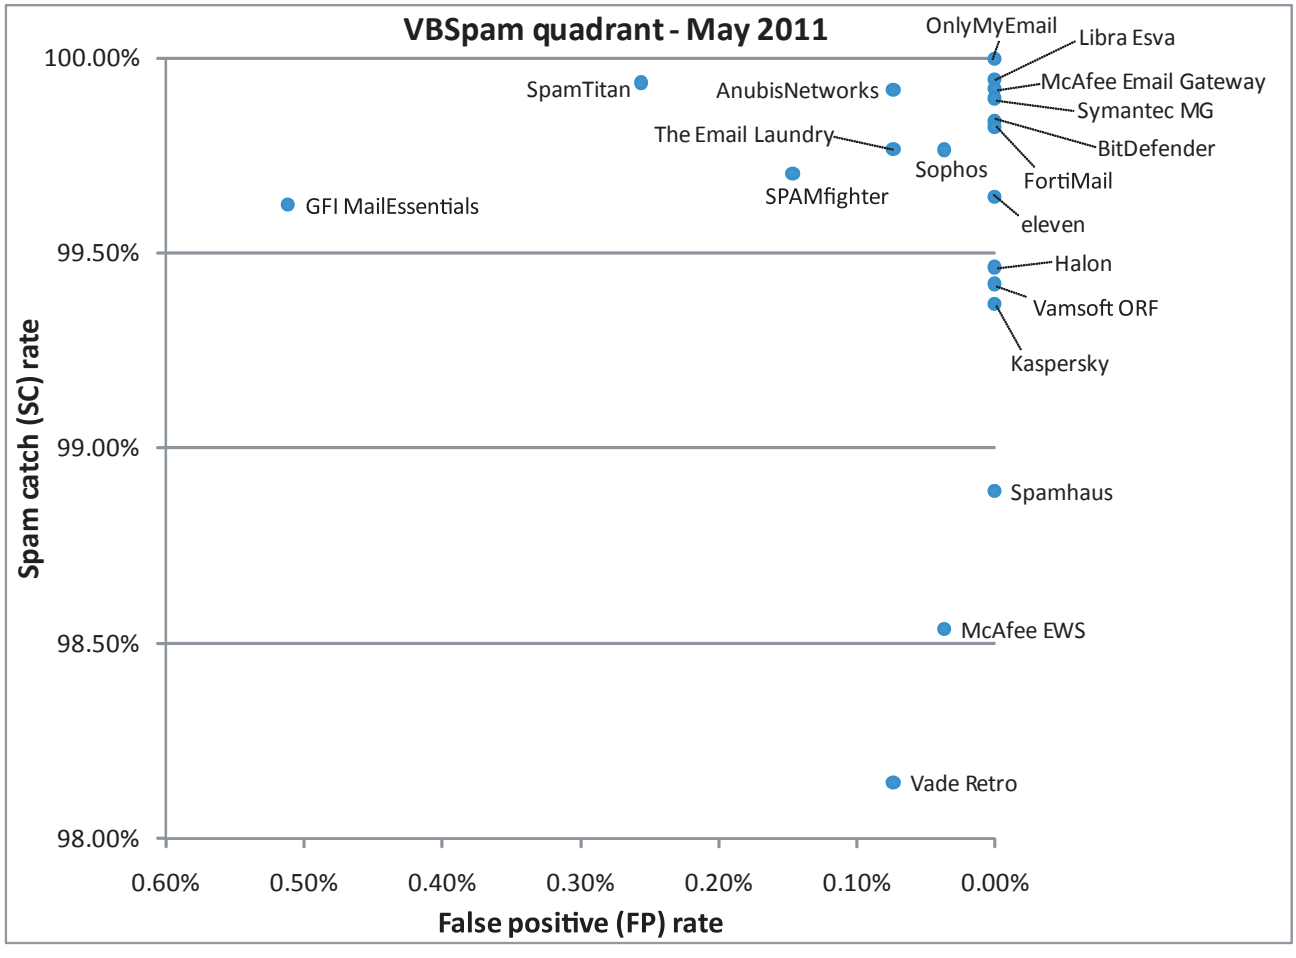

(Please refer to text for full product names.)

This month's test marks *SpamTitan*'s tenth entry and also the product's tenth VBSpam award. It had been some time since I last looked at the interface, but on using it to add black- and whitelists I remembered how easily it worked. Whitelisting by IP, however, is not available.

As in previous tests, *SpamTitan* blocked close to 100 per cent of all spam, though the second highest false positive rate in this month's test means there is something that can be improved upon.

Symantec Messaging Gateway 9.5 powered by Brightmail

SC rate: 99.90%

FP rate: 0.00%

Final score: 99.90

Project Honey Pot SC rate: 99.86%

Abusix SC rate: 99.94%

Symantec Messaging Gateway 9.5 powered by Brightmail (formerly *Symantec Brightmail Gateway*) is the newest version of the product, and we used the change in product version as an opportunity to move it to a *VMware ESXi 4.1* virtual server. The 9.5 version includes what the developers believe to be improvements to the anti-spam engine and I was curious to see if this would be reflected in the product's performance. It was: the spam catch rate – which was already excellent – was improved upon, and there were no false positives this time. The product thus wins its ninth VBSpam award in as many tests with the fourth highest final score.

What has not changed – at least not in a noticeable way – is the web interface to control the (virtual) appliance, and I was rather pleased by that. Adding IP and domain black- and whitelists and telling the product to do specific things when these lists were triggered was easy and worked as expected.

Products ranked by final score	Final score
OnlyMyEmail	100.00
Libra Esva	99.94
McAfee Email Gateway	99.92
Symantec Brightmail Gateway	99.90
BitDefender	99.84
FortiMail	99.82
eleven	99.64
Sophos Email Appliance	99.58
AnubisNetworks	99.55
Halon Security	99.46
ORF	99.42
The Email Laundry	99.40
Kaspersky Anti-Spam	99.37
SPAMfighter	98.97
Spamhaus	98.89
SpamTitan	98.66
McAfee EWS	98.35
Vade Retro	97.78
GFI MailEssentials	97.07

The Email Laundry

SC rate: 99.77%
FP rate: 0.07%
Final score: 99.40
Project Honey Pot SC rate: 99.67%
Abusix SC rate: 99.89%
SC rate pre-DATA: 99.30%

The Email Laundry does not allow customers to black- or whitelist by IP address and it should be noted once again that we do not wish to make assertions about whether or not this is a good thing; it is certainly something where human mistakes can have huge consequences. Domain black- and whitelisting is possible though, and was easily added in the product's interface.

As before, the product caught a large amount of spam, the vast majority of which was blocked at the SMTP level. As both the spam catch rate and the false positive rate improved, so did the final score and the product easily won its seventh VBSpam award.

Vade Retro Center

SC rate: 98.14%
FP rate: 0.07%
Final score: 97.78
Project Honey Pot SC rate: 99.13%
Abusix SC rate: 96.86%

Vade Retro offers a wide range of solutions, from hardware and virtual appliances to a number of hosted solutions. Unlike most of the other solutions, the hosted solution we have been testing does not allow for IP or domain black- and whitelisting.

The product scored well enough to achieve its seventh VBSpam award, but its developers may want to look at improving its spam catch rate – which was lower than any other product in this test. The fact that this was largely due to problems with the *Abusix* corpus may help them find the reason for this drop in performance.

Vamsoft ORF

SC rate: 99.42%
FP rate: 0.00%
Final score: 99.42
Project Honey Pot SC rate: 99.21%
Abusix SC rate: 99.69%

I have sung the praises of *ORF*'s user interface before, and using it to add black- and whitelists was once again a pleasure. Given the fact that *ORF* had no false positives for the fifth time in seven tests, though, few people are likely to need the whitelisting options.

While keeping the false positives to zero, *ORF* also managed to improve its spam catch rate, which was higher than in any previous test. A seventh VBSpam award will be proudly received at the company's Hungarian headquarters.

Spamhaus ZEN+DBL

SC rate: 98.89%
FP rate: 0.00%

Spamhaus ZEN+DBL contd.

Final score: 98.89
Project Honey Pot SC rate: 98.63%
Abusix SC rate: 99.23%
SC rate per-DATA: 98.15%

The increasing occurrence of URL shorteners in spam messages has presented *Spamhaus*'s *DBL* blocklist with a problem: blocking them would give false positives on the legitimate use of such shorteners; allowing them would give spammers a way to include their malicious URLs while avoiding detection.

The blocklist's developers have come up with a rather neat solution, returning different codes for known shorteners. This prevents messages containing them from being marked as spam, but allows users of the blacklists to resolve the real URL and hold this against the blacklist. (We did not do this in our tests.)

This ability to constantly adapt to spammers' techniques has earned *Spamhaus* eight VBSpam awards already and it adds a ninth to its tally in this test, catching a higher percentage of spam than on any previous occasion.

(As *Spamhaus* is only a partial solution, which needs to be integrated into a full solution, it does not make sense to black- or whitelist within this product.)

CONCLUSION

After two years of testing, and VBSpam certifications being awarded to 27 different products, readers should by now have a good picture of which products provide decent inbox protection and, more importantly, which of those can provide protection reliably over a prolonged period. We will, of course, continue to test products and award VBSpam certifications, but we also intend to provide more information about the products we look at.

The black- and whitelisting tests we introduced this month took some time to set up, but I think it was worth doing: with so many products performing so well, customers might want to differentiate between products by looking into the availability (or otherwise) of certain extras. We intend to look into more of these additional features in future tests.

The next VBSpam test will run in June 2011, with the results scheduled for publication in July. Developers interested in submitting products should email martijn.grooten@virusbtn.com.