

VBSPAM EMAIL SECURITY COMPARATIVE REVIEW MARCH 2026

Ionuț Răileanu & Adrian Luca

In the Q1 2026 VBSpam test – which forms part of *Virus Bulletin’s* continuously running security product test suite – we measured the performance of a number of email security solutions against various streams of wanted, unwanted and malicious emails. Half of the solutions we tested opted to be included in the public test, the rest opting for private testing (all details and results remaining unpublished). The solutions tested publicly – and included in this report – were nine full email security solutions and one open-source solution.

Our latest round of testing again revealed a diverse set of socially engineered and infrastructure-assisted email threats. Among the more notable examples were Finnish tax authority lures sent via abused *Freshservice* workflows,

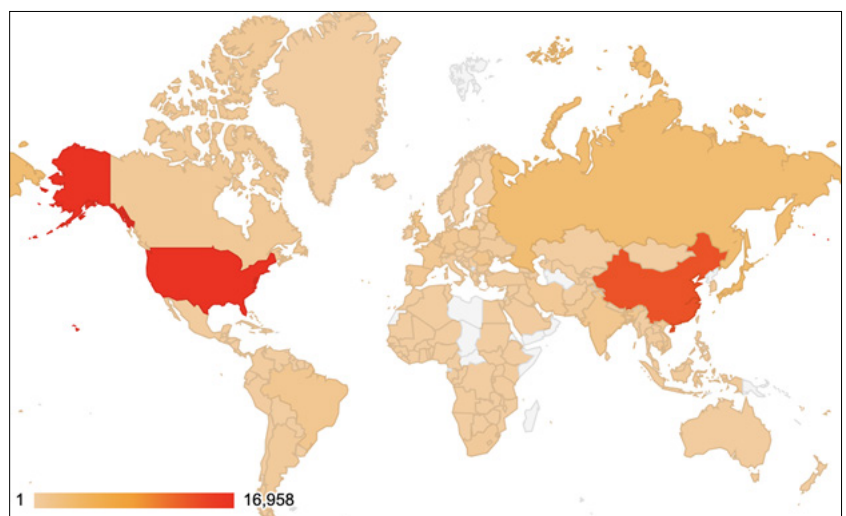
payment-themed brand impersonation delivered through *Atlassian Jira* notifications, and malicious content injected into an otherwise legitimate *Mailjet* newsletter campaign. These demonstrate how attackers continue to take advantage of trusted third-party delivery platforms, valid authentication and convincing branding in order to evade detection. Nevertheless, the stronger products in the test maintained very high spam, malware and phishing catch rates, with several solutions combining near-perfect detection with no false positives, although the gap to the weaker-performing solutions remained significant.

For some additional background to this report, the table and map below show the geographical distribution (based on sender IP address) of the spam emails seen in the test¹. (*Note: these statistics are relevant only to the spam samples we received during the test period.*)

¹ For 4,684 spam samples (7.42% of the total) we were unable to find data about geographical location based on IP address.

#	Sender’s IP country	Percentage of spam
1	United States	26.86%
2	China	21.00%
3	Japan	5.23%
4	Russian Federation	4.68%
5	Romania	2.04%
6	United Kingdom	1.90%
7	India	1.74%
8	Brazil	1.65%
9	France	1.37%
10	Germany	1.33%

Top 10 countries from which spam was sent.



Geographical distribution of spam based on sender IP address.

AMTSO STANDARD COMPLIANCE

This test was executed in accordance with the AMTSO Standard of the Anti-Malware Testing Standards Organization. The compliance status can be verified on the AMTSO website.

- **AMTSO Test ID:** AMTSO-LS1-TP179
- **Link:** <https://www.amtso.org/tests/virus-bulletin-vbspam-q1-2026/>

HIGHLIGHTS

Finnish tax authority impersonation via ticket notification

This is a high-risk phishing message delivered through a legitimate *Freshservice* notification channel, allowing it to bypass standard email security controls. The attacker embedded Finnish-language content in a ticket notification, impersonating *Verohallinto* and the *My Suomi* service and claiming that an official message was waiting for the recipient, possibly related to a tax refund or a request for additional information.

The message contains a malicious `links[.]truthsocial[.]com` URL behind a ‘Siirry viestiin’ (‘Go to message’) button, while using legitimate *Freshservice* infrastructure and *Suomi.fi* branding assets to increase credibility. The combination of trusted delivery, government impersonation, and a tax-themed lure represents a significant social engineering threat.

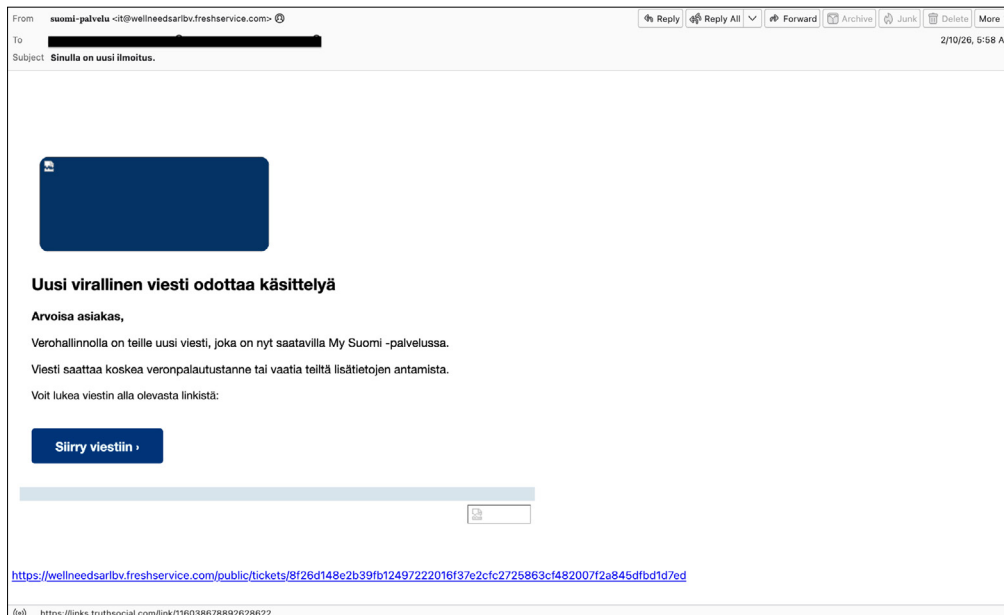
Brand impersonation via Jira ticket notification

Another phishing campaign that caught our attention targeted Danish users with a fake *Saphe* payment update notice delivered through a legitimate-looking *Atlassian Jira Service Management* notification flow. We observed near-identical messages in a short bursts from 1 to 3 February 2026, all using the subject ‘Midlertidigt betalingsproblem’ followed by a unique ticket ID and presenting polished *Saphe* branding, Danish payment-themed wording and a call to ‘Opdater betalingsoplysninger’ (‘Update payment information’).

The campaign is notable because the messages were DKIM-aligned with `atlassian.net` and sent via *Amazon SES*, giving them the appearance of legitimate service-desk traffic, while the HTML call-to-action redirected users to an unrelated `links[.]truthsocial[.]com` URL. This combination of trusted infrastructure, brand impersonation and a payment-failure lure makes the campaign a high-risk credential- and payment-harvesting threat that could plausibly evade standard email filtering.

Newsletter content injection via legitimate Mailjet campaign

On 12 February, for around five minutes only (17:11 to 17:15) we noticed a campaign of high-risk malicious messages embedded inside what otherwise appeared to be a legitimate *Mailjet*-delivered newsletter from ‘Institut du Tourisme - OTB’. The sender, subject, campaign metadata



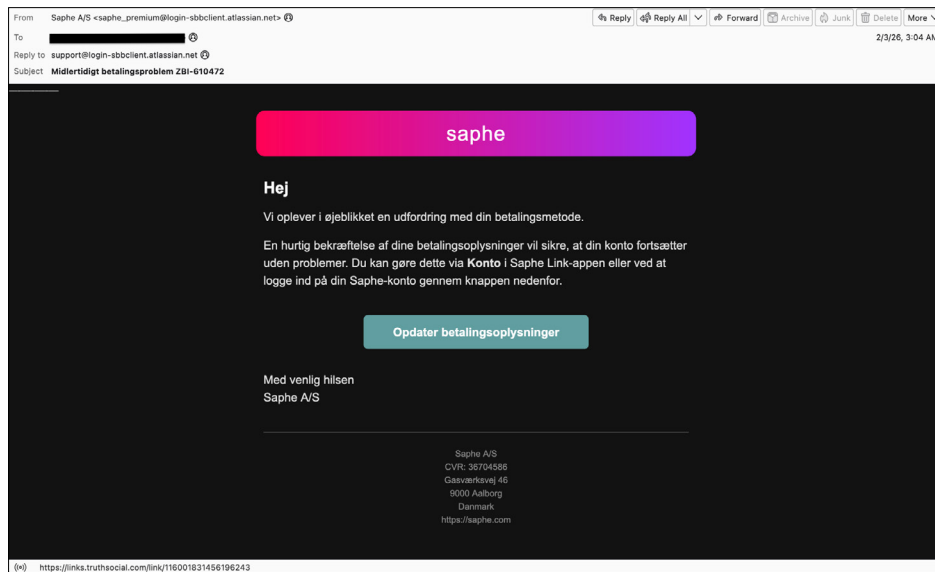
Finnish tax authority phishing sample.

and DKIM alignment with institut-tourisme[.]bzh were consistent across the observed samples, and most of the body of the email contained coherent content promoting French tourism training courses in 2026.

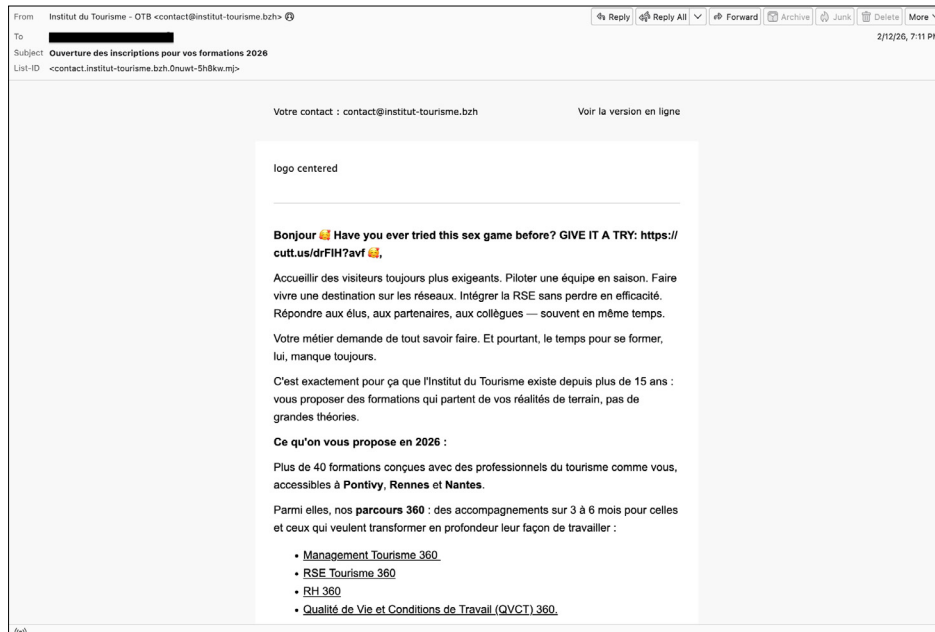
However, a small block near the top of the message had been altered to include fraudulent dating- or adult-themed lure texts such as ‘sent you a private message’, ‘want to meet you’ or ‘viewed your profile’, redirecting recipients to

a cutt[.]us/drFIH URL unrelated to the sender or the newsletter content.

This pattern suggests content injection into a legitimate mailing workflow, most likely through compromise of a *Mailjet* account, API key, connected CMS or reusable newsletter template, allowing an attacker to insert malicious content into an otherwise trustworthy campaign and increasing the likelihood of bypassing standard email security controls.



Phishing email sample impersonating Jira notification.



Newsletter content injection sample.

RESULTS

Of the nine participating full solutions, *Zoho Mail* achieved a VBSpam award, while six others were awarded VBSpam+ certification: *Bitdefender GravityZone Premium*, *FortiMail*, *N-able Mail Assure*, *N-able SpamExperts*, *Net at Work NoSpamProxy* and *SEPPmail.cloudfilter*.

(Note: since, for a number of products, catch rates and/or final scores were very close to, whilst remaining a fraction below, 100%, we quote all the spam-related scores with three decimal places.)

Bitdefender GravityZone Premium

SC rate: 99.979%
 FP rate: 0.00%
 Final score: 99.979
 Malware catch rate: 100.000%
 Phishing catch rate: 99.940%
 Project Honey Pot SC rate: 99.997%
 Abusix SC rate: 99.956%
 MXMailData SC rate: 100.000%
 Newsletters FP rate: 0.0%
 Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●



Another VBSpam+ certification is awarded to *Bitdefender*, which continues its long run of top-tier performances. With no false positives of any kind, a 100% malware catch rate and a 99.94% phishing catch rate, its results are among the strongest in the test.

Coro Email Security

SC rate: 92.759%
 FP rate: 0.00%
 Final score: 92.759
 Malware catch rate: 81.750%
 Phishing catch rate: 75.070%
 Project Honey Pot SC rate: 98.944%
 Abusix SC rate: 94.044%
 MXMailData SC rate: 56.050%
 Newsletters FP rate: 0.0%
 Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Coro Email Security showed a mixed performance in this test. While the product maintained zero false positives and delivered strong results on the *Project Honey Pot* and *Abusix* feeds, its malware, phishing and *MXMailData* catch rates were considerably lower than those of most other full solutions. Despite acceptable speed measurements, the

lower detection rates limited *Coro* to a final score below the threshold for VBSpam certification.

Fortinet FortiMail

SC rate: 99.892%
 FP rate: 0.00%
 Final score: 99.892
 Malware catch rate: 99.840%
 Phishing catch rate: 99.830%
 Project Honey Pot SC rate: 99.910%
 Abusix SC rate: 99.881%
 MXMailData SC rate: 99.860%
 Newsletters FP rate: 0.0%
 Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●



With no false positives of any kind and a spam catch rate close to 99.9%, *FortiMail* once again delivered a strong overall performance. Malware and phishing detection were both very high, and all speed measurements were green. These results comfortably earn *FortiMail* another VBSpam+ certification, with a final score of 99.892.

N-able Mail Assure

SC rate: 99.771%
 FP rate: 0.00%
 Final score: 99.771
 Malware catch rate: 100.000%
 Phishing catch rate: 99.850%
 Project Honey Pot SC rate: 99.965%
 Abusix SC rate: 99.527%
 MXMailData SC rate: 100.000%
 Newsletters FP rate: 0.0%
 Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●



N-able Mail Assure put in another strong and balanced performance, earning the product VBSpam+ certification. It blocked all malware samples, caught 99.85% of phishing emails, and maintained a 99.77% overall spam detection rate without generating any false positives.

N-able SpamExperts

SC rate: 99.773%
 FP rate: 0.00%
 Final score: 99.773
 Malware catch rate: 100.000%
 Phishing catch rate: 99.850%
 Project Honey Pot SC rate: 99.965%
 Abusix SC rate: 99.530%



MXMailData SC rate: 100.000%

Newsletters FP rate: 0.0%

Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Putting in a similarly impressive performance to that of its sister product, *N-able SpamExperts* also earns VBSpam+ certification, with a final score of 99.773.

Net at Work NoSpamProxy

SC rate: 99.929%

FP rate: 0.00%

Final score: 99.929

Malware catch rate: 100.000%

Phishing catch rate: 99.920%

Project Honey Pot SC rate: 99.990%

Abusix SC rate: 99.853%

MXMailData SC rate: 100.000%

Newsletters FP rate: 0.0%

Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

NoSpamProxy delivered another excellent performance in this test. With 100% malware protection, a 99.92% phishing catch rate, a 99.93% overall spam catch rate and no false positives, the product remains one of the best performers in the field. Green values on all speed measurements help secure a well deserved VBSpam+ certification with a final score of 99.929.



Rspamd

SC rate: 75.199%

FP rate: 0.47%

Final score: 72.848

Malware catch rate: 37.560%

Phishing catch rate: 78.240%

Project Honey Pot SC rate: 79.517%

Abusix SC rate: 75.006%

MXMailData SC rate: 54.860%

Newsletters FP rate: 0.0%

Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

The malware corpus was particularly challenging for the open-source *Rspamd* in this test, with only 37.56% of the malicious samples blocked. Although the product's phishing catch rate was noticeably better and its speed measurements remained green throughout, the combination of a low spam catch rate and a high false positive rate resulted in a final score falling below the threshold for certification. The result underlines the gap between the basic open-source configuration and the stronger full commercial solutions in this test.

Rspamd Premium 3.14.3

SC rate: 98.484%

FP rate: 0.47%

Final score: 96.133

Malware catch rate: 99.110%

Phishing catch rate: 98.360%

Project Honey Pot SC rate: 97.515%

Abusix SC rate: 99.283%

MXMailData SC rate: 99.380%

Newsletters FP rate: 0.0%

Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

The premium *Rspamd* configuration significantly outperformed the basic version, especially on malware, phishing and overall spam detection. Despite this, a false positive rate of 0.47% had a clear impact on the final score, which fell short of the certification threshold despite otherwise strong results.

SEPPmail.cloudfilter

SC rate: 99.978%

FP rate: 0.00%

Final score: 99.978

Malware catch rate: 100.000%

Phishing catch rate: 99.870%

Project Honey Pot SC rate: 99.969%

Abusix SC rate: 99.993%

MXMailData SC rate: 99.950%

Newsletters FP rate: 0.0%

Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

SEPPmail delivered another near-flawless performance in this test. With no false positives of any kind, a perfect malware catch rate, a 99.87% phishing catch rate and a 99.978% overall spam detection rate, it ranked among the strongest solutions in the test. Green values for all speed measurements support an excellent final score of 99.978 and VBSpam+ certification.



Zoho Mail

SC rate: 98.025%

FP rate: 0.00%

Final score: 98.025

Malware catch rate: 100.000%

Phishing catch rate: 99.080%

Project Honey Pot SC rate: 98.341%

Abusix SC rate: 97.306%

MXMailData SC rate: 99.950%



Newsletters FP rate: 0.0%

Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Zoho Mail achieved very strong results on the malware, phishing and overall spam corpora, while keeping false positives at zero. Although its catch rates were not quite as high as those of the very top performers, the product delivered a well balanced result with green values for all speed measurements. With a final score of 98.025, *Zoho Mail* earns VBSspam certification.

APPENDIX: SET-UP, METHODOLOGY AND EMAIL CORPORA

The full VBSspam test methodology can be found at <https://www.virusbulletin.com/testing/vbspam/vbspam-methodology/vbspam-methodology-ver30/>.

The test ran for 15 days, from 12am on 1 February to 12am on 16 February 2026 (GMT).

The test corpus consisted of 64,016 emails. 63,143 of these were spam, 28,833 of which were provided by *Project Honey Pot*, 28,462 were provided by *Abusix*, with the remaining 5,848 spam emails provided by *MXMailData*. There were 845 legitimate emails ('ham') and 28 newsletters, a category that includes various kinds of commercial and non-commercial opt-in mailings.

161 emails in the spam corpus were considered 'unwanted' (see the June 2018 report²) and were included with a weight of 0.2; this explains the non-integer numbers in some of the tables.

Moreover, 1,901 emails from the spam corpus were found to contain a malicious attachment while 4,762 contained a link to a phishing or malware site; though we report separate performance metrics on these corpora, it should be noted that these emails were also counted as part of the spam corpus.

Emails were sent to the products in real time and in parallel. Though products received the email from a fixed IP address, all products had been set up to read the original sender's IP address as well as the EHLO/HELO domain sent during the SMTP transaction, either from the email headers or through an optional XCLIENT SMTP command³.

For those products running in our lab, we all ran them as virtual machines on a *VMware ESXi* cluster. 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

² <https://www.virusbulletin.com/virusbulletin/2018/06/vbspam-comparative-review>

³ http://www.postfix.org/XCLIENT_README.html

and note that the amount of email we receive is representative of that received by a small organization.

Although we stress that different customers have different needs and priorities, and thus different preferences when it comes to the ideal ratio of false positive to false negatives, we created a one-dimensional 'Final score' to compare products. This 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:

$$\text{WFP rate} = (\# \text{false positives} + 0.2 * \min(\# \text{newsletter false positives}, 0.2 * \# \text{newsletters})) / (\# \text{ham} + 0.2 * \# \text{newsletters})$$

while in the spam catch rate (SC), emails considered 'unwanted' (see above) are included with a weight of 0.2. The Final score is then defined as:

$$\text{Final score} = \text{SC} - (5 \times \text{WFP})$$

In addition, for each product, we measure how long it takes to deliver emails from the ham corpus (excluding false positives) and, after ordering these emails by this time, we colour-code the emails at the 10th, 50th, 95th and 98th percentiles:

- (green) = up to 30 seconds
- (yellow) = 30 seconds to two minutes
- (orange) = two to ten minutes
- (red) = more than ten minutes

Products earn VBSspam certification if the value of the final score is at least 98 and the 'delivery speed colours' at 10 and 50 per cent are green or yellow and that at 95 per cent is green, yellow or orange.








Meanwhile, products that combine a spam catch rate of 99.5% or higher with a lack of false positives, no more than 2.5% false positives among the newsletters and 'delivery speed colours' of green at 10 and 50 per cent and green or yellow at 95 and 98 per cent earn a VBSspam+ award.

Head of Testing: Peter Karsai
 Security Test Engineers: Klaudia Kittı Csia, Adrian Luca, Ionuț Răileanu
 Senior Threat Analyst: Norbert Biro
 Operations Manager: Bálint Tanos
 Sales Executive: Allison Sketchley
 Marketing Manager: David Kelemen
 Editorial Assistant: Helen Martin

© 2025 Virus Bulletin Ltd, Manor House, Howbery Business Park, Wallingford OX10 8BA, UK

Tel: +44 20 3920 6348 Email: editorial@virusbulletin.com

Web: <https://www.virusbulletin.com/>

	True negatives	False positives	FP rate	False negatives	True positives	SC rate	Final score	VBSpam
Bitdefender GravityZone Premium	845	0	0.00%	13.4	63000.8	99.979%	99.979	
Coro Email Security	845	0	0.00%	4562.8	58451.4	92.759%	92.759	
Fortinet FortiMail	845	0	0.00%	67.8	62946.4	99.892%	99.892	
N-able Mail Assure	845	0	0.00%	144.2	62870	99.771%	99.771	
N-able SpamExperts	845	0	0.00%	143.2	62871	99.773%	99.773	
Net at Work NoSpamProxy	845	0	0.00%	44.6	62969.6	99.929%	99.929	
Rspamd	841	4	0.47%	15628.2	47386	75.199%	72.848	
Rspamd Premium 3.14.3	841	4	0.47%	955.4	62058.8	98.484%	96.133	
SEPPmail.cloudfilter	845	0	0.00%	14	63000.2	99.978%	99.978	
Zoho Mail	845	0	0.00%	1244.8	61769.4	98.025%	98.025	

	Newsletters		Malware		Phishing		Project Honey Pot		Abusix		MXMailData		STDev†
	False positives	FP rate	False negatives	SC rate	False negatives	SC rate	False negatives	SC rate	False negatives	SC rate	False negatives	SC rate	
Bitdefender GravityZone Premium	0	0.00%	0	100.0000%	3	99.940%	1	99.997%	12.4	99.956%	0	100.0000%	0.12
Coro Email Security	0	0.00%	347	81.750%	1187	75.070%	304.2	98.944%	1688.6	94.044%	2570	56.050%	9.38
Fortinet FortiMail	0	0.00%	3	99.840%	8	99.830%	26	99.910%	33.8	99.881%	8	99.860%	0.36
N-able Mail Assure	0	0.00%	0	100.0000%	7	99.850%	10	99.965%	134.2	99.527%	0	100.0000%	0.65
N-able SpamExperts	0	0.00%	0	100.0000%	7	99.850%	10	99.965%	133.2	99.530%	0	100.0000%	0.65
Net at Work NoSpamProxy	0	0.00%	0	100.0000%	4	99.920%	3	99.990%	41.6	99.853%	0	100.0000%	0.2
Rspamd	0	0.00%	1187	37.560%	1036	78.240%	5902.6	79.517%	7085.6	75.006%	2640	54.860%	7.9
Rspamd Premium 3.14.3	0	0.00%	17	99.110%	78	98.360%	716.2	97.515%	203.2	99.283%	36	99.380%	1.71
SEPPmail. cloudfilter	0	0.00%	0	100.0000%	6	99.870%	9	99.969%	2	99.993%	3	99.950%	0.22
Zoho Mail	0	0.00%	0	100.0000%	44	99.080%	478	98.341%	763.8	97.306%	3	99.950%	3.54

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

	Speed			
	10%	50%	95%	98%
Bitdefender GravityZone Premium	●	●	●	●
Coro Email security	●	●	●	●
Fortinet FortiMail	●	●	●	●
N-able Mail Assure	●	●	●	●
N-able SpamExperts	●	●	●	●
Net at Work NoSpamProxy	●	●	●	●
Rspamd	●	●	●	●
Rspamd Premium 3.14.3	●	●	●	●
SEPPmail.cloudfilter	●	●	●	●
Zoho Mail	●	●	●	●

● 0–30 seconds; ● 30 seconds to two minutes; ● two minutes to 10 minutes; ● more than 10 minutes.

Products ranked by final score	
Bitdefender GravityZone Premium	99.979
SEPPmail.cloudfilter	99.978
Net At Work NoSpamProxy	99.929
Fortinet FortiMail	99.892
N-able SpamExperts	99.773
N-able Mail Assure	99.771
Zoho Mail	98.025
Rspamd Premium 3.14.3	96.133
Coro Email Security	92.759
Rspamd	72.848

Hosted solutions	Anti-malware	IPv6	DKIM	SPF	DMARC	Multiple MX-records	Multiple locations
Coro Email security	Coro	√	√	√	√		
N-able Mail Assure	N-able Mail Assure	√	√	√	√		
N-able SpamExperts	SpamExperts	√	√	√	√		
Net at Work NoSpamProxy	32Guards & NoSpamProxy		√	√	√	√	√
Rspamd Premium	ClamAV		√	√	√	√	√
SEPPmail.cloudfilter	SEPPmail, ClamAV & ESET	√	√	√	√	√	√
Zoho Mail	Zoho		√	√	√	√	√

Local solutions	Anti-malware	IPv6	DKIM	SPF	DMARC	Interface			
						CLI	GUI	Web GUI	API
Bitdefender GravityZone Premium	Bitdefender	√				√		√	√
Fortinet FortiMail	Fortinet	√	√	√	√	√		√	√
Rspamd	None					√			

