DECEMBER 2025

virus BULLETIN

Covering the global threat landscape

VBSPAM EMAIL SECURITY COMPARATIVE REVIEW DECEMBER 2025

Ionuţ Răileanu & Adrian Luca

In the Q4 2025 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tested opted to be included in the public test, the rest opting for private testing (all details and results remaining unpublished).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 ten full email security solutions and one open-source solution.

Our latest round of testing once again revealed some sophisticated and targeted email threats, but we also observed continued adaptation and overall improvement in the filtering capabilities of email security solutions.

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 test[1]. (*Note: these statistics are relevant only to the spam samples we received during the test period.*)

AMTSO STANDARD COMPLIANCE

This test was executed in accordance with the AMTSO Standard of the Anti-Malware Testing Standards Organization.

[1] For a number of samples (6,561 samples; 6.50% of the total) we were unable to find data about geographical location based on IP address.

#	Sender's IP country	Percentage of spam
1	China	34.56%
2	United States	17.81%
3	Japan	4.59%
4	Brazil	3.13%
5	Russian Federation	3.09%
6	India	1.91%
7	Argentina	1.80%
8	Germany	1.47%
9	United Kingdom	1.22%
10	France	1.09%

Top 10 countries from which spam was sent.

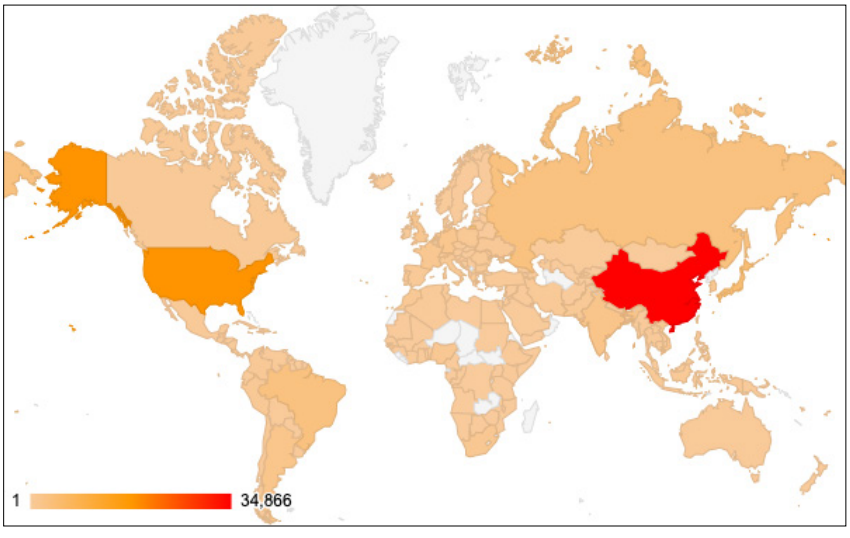

Geographical distribution of spam based on sender IP address.

ISSN 1749-7027

The compliance status can be verified on the AMTSO website:

- **AMTSO Test ID**: AMTSO-LS1-TP166
- **Link**: https://www.amtso.org/tests/virus-bulletin-vbspam-q4-2025/

HIGHLIGHTS

Office 365 phishing campaign

A phishing campaign that caught our attention was one that targeted users of *Office 365*. We observed the active campaign on 3 November from 15:12 to 17:33 (UTC), bypassing the filters of most of the email security solutions.

The campaign likely aims to harvest *Microsoft 365* credentials or collect payment data via phone. The HTML attachment is engineered to load a live phishing page; if opened in a browser with network access, it can capture credentials.

Indicators of compromise (IOCs)

- **Subject**: Re: Your Office 365=C2=AE Subscription Has Been Renewed
- **From**: info@se[.]df[.]gov[.]br (spoofed)
- **Return-Path**: LNNRUWL2DTQCVN@se[.]df[.]gov[.]br
- **Source IP**: 185[.]236[.]231[.]236 (SPF fail)
- **Message-Id domain**: repfitness[.]com (mismatch)
- **Attachment**: Microsoft Subscription Team_[0-9]{7}[.]htm – HTML with <iframe> to external site on port 2087.
- **URL** (from the HTML attachment): hxxps://kicvv7ru7i44weejzze075fip2wy4t[.]eliteretouching[.]com:2087/impact?...=user
- **Vishing lure**: Phone number: +1-720-420-1731

Brand impersonation via ticket notification

This email is a high-risk phishing message delivered through a legitimate *Freshservice* notification channel, allowing it to bypass standard email security controls. The attacker inserted fraudulent content into a ticket request, impersonating *AG Insurance* and claiming the recipient is entitled to a reimbursement.

The message contains a malicious *Thinkific* tracking link that redirects to a phishing site (ag[.]assurance-contact[.]info) designed to harvest personal and financial information. The email uses professional formatting, brand impersonation, and urgency to appear credible, representing a significant social-engineering threat.

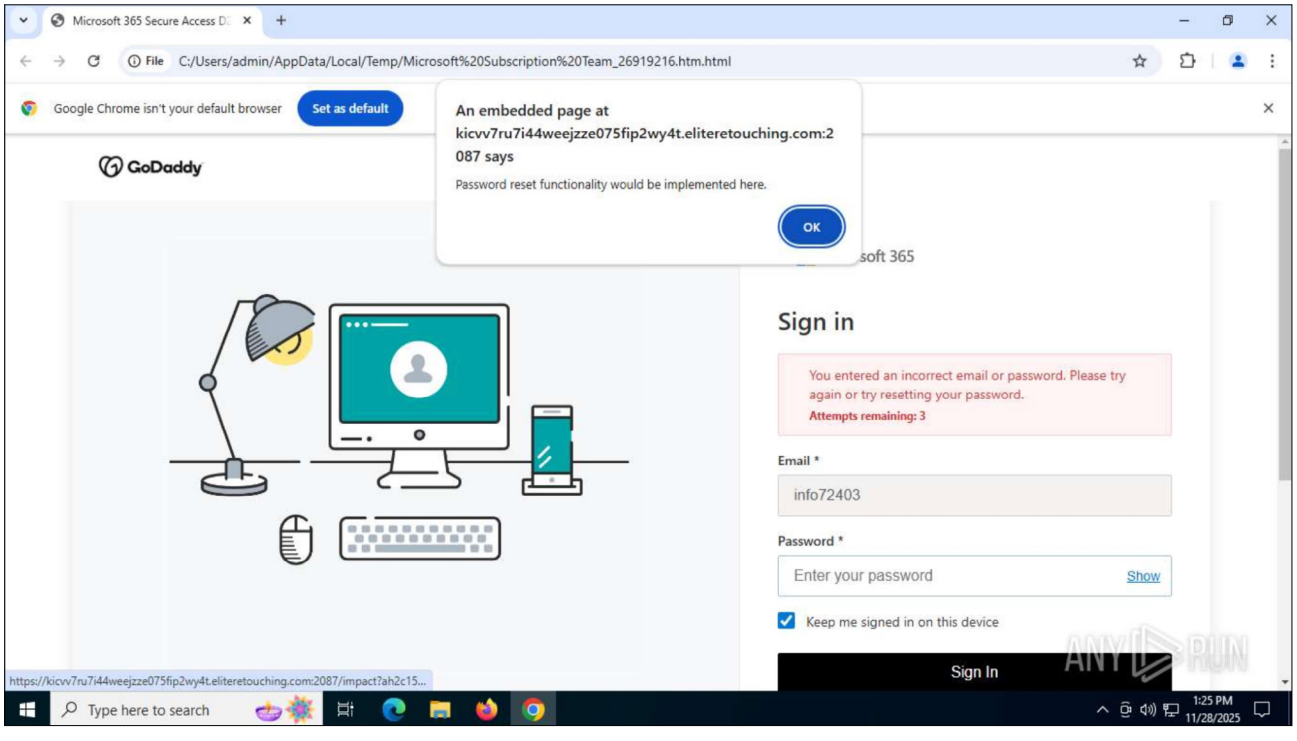

Phishing page loaded from the HTM attachment.

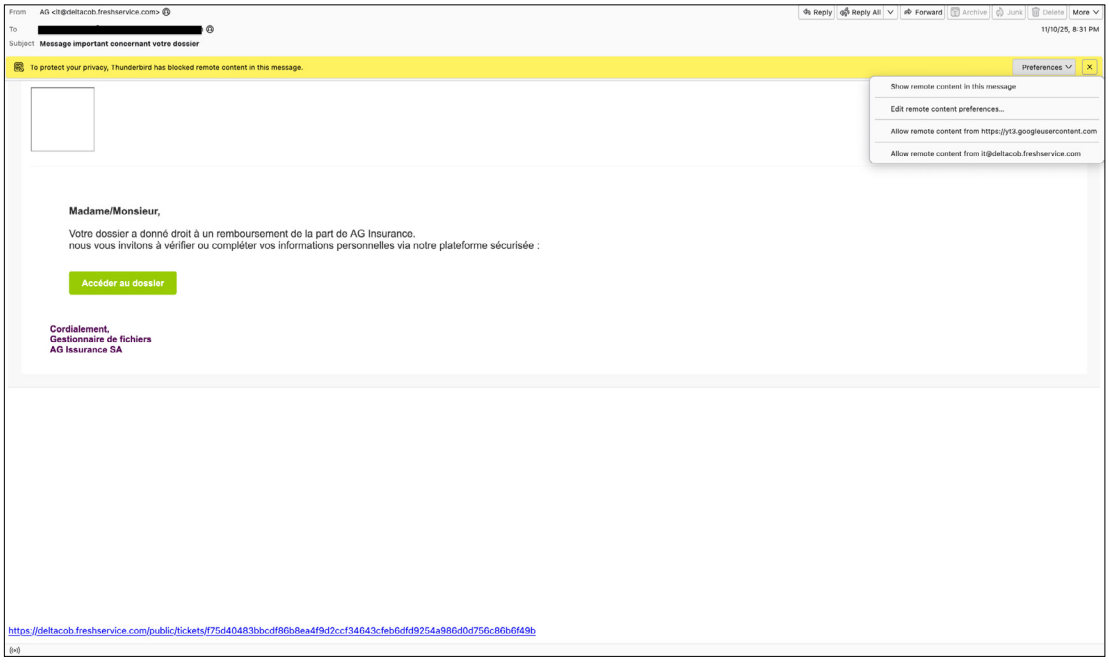

Brand impersonation via ticket notification sample.

Indicators of compromise (IOCs)

Malicious domains / URLs:

- ag[.]assurance-contact[.]info – final phishing landing domain.
- thinkific[.]com / api[.]thinkific[.]com – used as a redirect/tracking layer to mask the phishing link.
- filesusr[.]com (*Thinkific* asset host) – used as part of the redirection chain.

Suspicious email artifacts:

- Compromised *Freshservice* ticket system – phishing content embedded inside a legitimate service notification.
- Spoofed sender identity – *AG Insurance* impersonation within the ticket message body.
- Embedded phishing call to action – 'Follow-up here' link pointing to malicious redirect chain.

Phishing infrastructure characteristics:

- Newly registered / low-reputation phishing domain (*[.]assurance-contact[.]info).
- External hosting with no alignment to the impersonated brand (*AG Insurance*).
- Use of multiple redirects to evade detection and reputation filtering.

RESULTS

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

(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: 99.992%
FP rate: 0.00%
Final score: 99.992
Malware catch rate: 100.000%
Phishing catch rate: 100.000%
Project Honey Pot SC rate: 99.996%
Abusix SC rate: 99.987%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Another VBSpam+ certification is awarded to *Bitdefender*, marking an uninterrupted 10-year run of the higher level certification. The product's performance continues to be among the top ranking, with no malware or phishing threats bypassing *Bitdefender*'s filters and no false positives of any kind.

Coro Email security

SC rate: 99.758%
FP rate: 0.38%
Final score: 97.884
Malware catch rate: 100.000%
Phishing catch rate: 99.910%
Project Honey Pot SC rate: 99.696%
Abusix SC rate: 99.817%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

On its first appearance in the VBSpam test *Coro Email security* performed impressively, blocking more than 99.91% of phishing samples and 100% of malware threats, with an overall 99.75% spam detection rate. Four false positives were the only fly in the ointment, resulting in the final score falling just short of the VBSpam award criteria.

Fortinet FortiMail

SC rate: 99.920%
FP rate: 0.00%
Final score: 99.920
Malware catch rate: 100.000%
Phishing catch rate: 99.860%
Project Honey Pot SC rate: 99.896%
Abusix SC rate: 99.946%
MXMailData SC rate: 99.970%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

With no false positives of any kind, higher than 99.90% spam catch rate and green values for all speed measurements, *FortiMail* continues the same solid performance we've seen from it in previous tests and earns VBSpam+ certification.

N-able Mail Assure

SC rate: 99.938%
FP rate: 0.00%
Final score: 99.938
Malware catch rate: 100.000%
Phishing catch rate: 99.980%
Project Honey Pot SC rate: 99.983%
Abusix SC rate: 99.876%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

N-able Mail Assure earns VBSpam+ certification with an excellent and comprehensive performance. It boasts a 99.93% spam detection rate and zero false positives, leading to a final score of 99.938.

N-able SpamExperts

SC rate: 99.936%
FP rate: 0.00%
Final score: 99.936
Malware catch rate: 100.000%
Phishing catch rate: 99.980%
Project Honey Pot SC rate: 99.983%
Abusix SC rate: 99.871%
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.

Net at Work NoSpamProxy

SC rate: 99.984%
FP rate: 0.00%
Final score: 99.984
Malware catch rate: 100.000%
Phishing catch rate: 100.000%
Project Honey Pot SC rate: 99.997%
Abusix SC rate: 99.966%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

NoSpamProxy was one of only three products in this test that managed to successfully block all the malware and phishing threats. Adding a 99.98% spam catch rate, a lack of false positives and green values for all speed measurements, a well deserved VBSpam+ certification is awarded to *Net at Work*'s product.

Rspamd

SC rate: 79.163%
FP rate: 0.28%
Final score: 77.663
Malware catch rate: 82.300%
Phishing catch rate: 81.660%
Project Honey Pot SC rate: 74.690%
Abusix SC rate: 85.284%
MXMailData SC rate: 71.640%
Newsletters FP rate: 3.2%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

As in previous tests, the open-source *Rspamd* found dealing with the malware and phishing samples a challenge and in this test achieved a slightly lower final score than previously. Despite this, we continue to see a decent overall performance from the solution.

Rspamd Premium 3.13.2

SC rate: 98.770%
FP rate: 0.57%
Final score: 95.959
Malware catch rate: 99.830%
Phishing catch rate: 99.240%
Project Honey Pot SC rate: 98.299%
Abusix SC rate: 99.579%
MXMailData SC rate: 95.800%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

The upgraded *Rspamd* configuration significantly outperformed the basic version, successfully blocking 98.77% of spam samples and achieving a final score of 95.959.

SEPPmail.cloudfilter

SC rate: 99.9998%
FP rate: 0.00%
Final score: 99.9998
Malware catch rate: 100.000%
Phishing catch rate: 100.000%
Project Honey Pot SC rate: 100.000%
Abusix SC rate: 99.9998%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

It was an almost flawless performance from *SEPPmail* in this test. With no false positives of any kind and only one missed unwanted sample, the product earns VBSpam+ certification with an almost perfect final score.

Sophos Email

SC rate: 99.926%
FP rate: 0.00%
Final score: 99.926
Malware catch rate: 99.140%
Phishing catch rate: 99.990%
Project Honey Pot SC rate: 99.920%
Abusix SC rate: 99.950%
MXMailData SC rate: 99.690%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

In this test, *Sophos* managed to block more than 99.90% of malware and phishing samples while successfully filtering the ham and newsletters with no false positives. With a final score of 99.926 and green values for all speed measurements, the product earns VBSpam+ certification.

Zoho Mail

SC rate: 99.119%
FP rate: 0.09%
Final score: 98.651
Malware catch rate: 99.830%
Phishing catch rate: 99.490%
Project Honey Pot SC rate: 99.034%
Abusix SC rate: 99.176%
MXMailData SC rate: 99.790%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Zoho Mail achieved higher than 99% catch rates across the malware and phishing samples as well as on the overall spam corpus. Whilst the product successfully filtered the newsletters with no false positives, a false positive in the ham set brought its final score beneath the VBSpam+ threshold – nevertheless, the product earns VBSpam certification with ease.

APPENDIX: SET-UP, METHODOLOGY AND EMAIL CORPORA

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

The test ran for 16 days, from 12am on 1 November to 12am on 17 November 2025 (GMT).

The test corpus consisted of 101,970 emails. 100,878 of these were spam, 54,084 of which were provided by *Project Honey Pot*, 43,532 were provided by *Abusix* with the remaining 3,262 spam emails provided by *MXMailData*. There were 1,061 legitimate emails ('ham') and 31 newsletters, a category that includes various kinds of commercial and non-commercial opt-in mailings.

16 emails in the spam corpus were considered 'unwanted' (see the June 2018 report[2]) and were included with a weight of 0.2; this explains the non-integer numbers in some of the tables.

Moreover, 582 emails from the spam corpus were found to contain a malicious attachment while 12,884 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 command[3].

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 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:

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

while in the spam catch rate (SC), emails considered 'unwanted' (see above) are included with a weight of 0.2.

[2] https://www.virusbulletin.com/virusbulletin/2018/06/vbspam-comparative-review

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

The Final score is then defined as:

Final score = SC - (5 x 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 VBSpam 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 VBSpam+ award.

Head of Testing: Peter Karsai
Security Test Engineers: Klaudia Kitti 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

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	1061	0	0.00%	7.8	100857.4	99.992%	99.992	SPAM + Verified
Coro Email security	1057	4	0.38%	244.2	100621	99.758%	97.884	
Fortinet FortiMail	1061	0	0.00%	80.4	100784.8	99.920%	99.920	SPAM + Verified
N-able Mail Assure	1061	0	0.00%	63	100802.2	99.938%	99.938	SPAM + Verified
N-able SpamExperts	1061	0	0.00%	65	100800.2	99.936%	99.936	SPAM + Verified
Net at Work NoSpamProxy	1061	0	0.00%	16	100849.2	99.984%	99.984	SPAM + Verified
Rspamd	1058	3	0.28%	21017.4	79847.8	79.163%	77.663	
Rspamd Premium 3.13.2	1055	6	0.57%	1240.2	99625	98.770%	95.959	
SEPPmail.cloudfilter	1061	0	0.00%	0.2	100865	99.9998%	99.9998	SPAM + Verified
Sophos Email	1061	0	0.00%	74.8	100790.4	99.926%	99.926	SPAM + Verified
Zoho Mail	1060	1	0.09%	888.2	99977	99.119%	98.651	SPAM Verified

	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	100.000%	0	100.000%	2	99.996%	5.8	99.987%	0	100.000%	0.11
Coro Email security	0	0.00%	0	100.000%	11	99.910%	164.6	99.696%	79.6	99.817%	0	100.000%	0.81
Fortinet FortiMail	0	0.00%	0	100.000%	18	99.860%	56	99.896%	23.4	99.946%	1	99.970%	0.32
N-able Mail Assure	0	0.00%	0	100.000%	2	99.980%	9	99.983%	54	99.876%	0	100.000%	0.21
N-able SpamExperts	0	0.00%	0	100.000%	2	99.980%	9	99.983%	56	99.871%	0	100.000%	0.22
Net at Work NoSpamProxy	0	0.00%	0	100.000%	0	100.000%	1.4	99.997%	14.6	99.966%	0	100.000%	0.13
Rspamd	1	3.23%	103	82.300%	2363	81.660%	13686.8	74.690%	6405.6	85.284%	925	71.640%	7.82
Rspamd Premium 3.13.2	0	0.00%	1	99.830%	98	99.240%	919.8	98.299%	183.4	99.579%	137	95.800%	1.29
SEPPmail.cloudfilter	0	0.00%	0	100.000%	0	100.000%	0	100.000%	0.2	99.9998%	0	100.000%	0.02
Sophos Email	0	0.00%	5	99.140%	1	99.990%	43	99.920%	21.8	99.950%	10	99.690%	0.35
Zoho Mail	0	0.00%	1	99.830%	66	99.490%	522.4	99.034%	358.8	99.176%	7	99.790%	2.01

† *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 3.13.2	🟢	🟢	🟢	🟢
SEPPmail.cloudfilter	🟢	🟢	🟢	🟢
Sophos Email	🟢	🟢	🟢	🟢
Zoho Mail	🟢	🟢	🟢	🟡

🟢 *0–30 seconds;* 🟡 *30 seconds to two minutes;* 🟠 *two minutes to 10 minutes;*
🔴 *more than 10 minutes.*

Products ranked by final score	
SEPPmail.cloudfilter	99.9998
Bitdefender GravityZone Premium	99.992
Net at Work NoSpamProxy	99.984
N-able Mail Assure	99.938
N-able SpamExperts	99.936
Sophos Email	99.926
Fortinet FortiMail	99.920
Zoho Mail	98.651
Coro Email security	97.884
Rspamd Premium 3.13.2	95.959

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	√	√	√	√	√	√
Sophos Email	Sophos	√	√	√	√	√	√
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					√