SEPTEMBER 2026

virus BULLETIN

Covering the global threat landscape

VBSPAM EMAIL SECURITY COMPARATIVE REVIEW SEPTEMBER 2026

Ionuţ Răileanu & Adrian Luca

In the Q3 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tested opted to be included in the public test, the rest opting for private testing (all details and results remaining unpublished).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 ten full email security solutions and one open-source solution.

The threats observed during this period demonstrate how modern phishing campaigns increasingly blend familiar business and consumer lures with lightweight evasive infrastructure. Rather than relying on obvious malware attachments, these messages used convincing account, invoice, banking, and subscription pretexts to push users toward external web flows where the real risk emerged only after redirects, browser checks, or unusual URL encoding. This made them difficult for email security solutions to block, as the malicious intent was distributed across social engineering, authenticated or plausible sender infrastructure, and cloaked destinations that could appear benign to static scanners or simplified sandbox analysis.

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.*)

[1] For 2,749 spam samples (2.66% 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	67.54%
2	China	7.06%
3	Russian Federation	3.36%
4	Japan	2.37%
5	Asia/Pacific Region	2.26%
6	United Kingdom	1.25%
7	Canada	0.99%
8	Germany	0.86%
9	Netherlands	0.84%
10	Brazil	0.74%

Top 10 countries from which spam was sent.

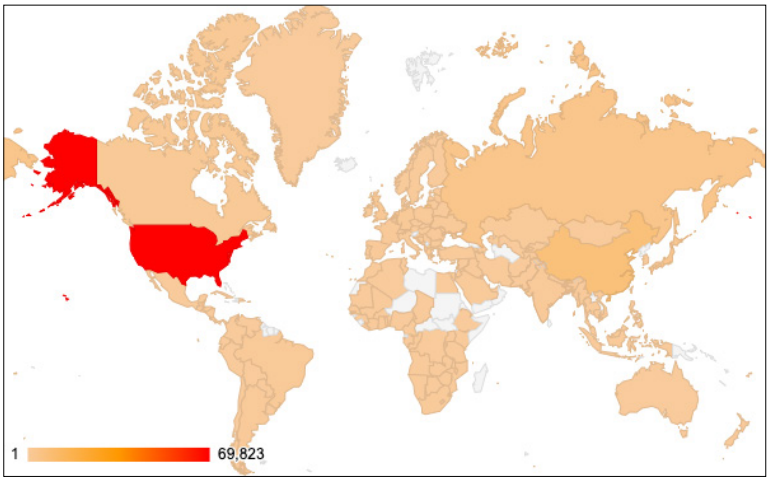

Geographical distribution of spam based on sender IP address.

ISSN 1749-7027

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-TP207
- **Link**: https://www.amtso.org/tests/virus-bulletin-vbspam-q3-2026/

HIGHLIGHTS

Antivirus renewal scareware phishing

During this period we observed a Dutch-language antivirus-themed phishing sample impersonating *McAfee*/*TotalAV* and warning recipients that their device was infected with '631 dangerous viruses'. The message used urgency, account-closure threats, and a 90% discount lure to push users towards a renewal action through sanitized redirect infrastructure such as `31-59-175-195[.]syd[.]nbn[.]aussiebb[.]net`, `loadswage[.]com`, and an unsubscribe path on `eightindigostove[.]com`.

The final threat was assessed as fake antivirus renewal scareware/subscription fraud, likely intended to monetize clicks through a deceptive renewal or affiliate funnel and potentially collect payment details, rather than deliver a confirmed malware payload.

The sample was challenging for security solutions because it contained no attachment or executable, relied on HTML-only social engineering, used a legitimate-looking sender domain, separated tracking/unsubscribe/CTA URLs, and routed users through live redirect infrastructure whose final destination could change, expire, or cloak by time and location.

Cloaked invoice phishing redirects to OpenSea

A notable August 2026 sample used a German overdue-payment-letter theme to disguise a Web3 fraud funnel as what appeared to be a routine accounting message. Sent through *Amazon SES* from the DKIM-aligned domain `moolaah[.]com`, the email referenced an unpaid invoice and invited the recipient to open a supposed 'Mahnschreiben' (overdue payment notice), although the message contained no attachment at all.

The call to action instead led to `website-2df62808[.]mvplineup[.]com/audacity/underside`, where the first-stage page served mostly decoy markup, hidden text, a zero-size iframe, and obfuscated JavaScript designed to collect browser and timezone signals before submitting a hidden POST request. Only after this gating step did the chain resolve to `opensea[.]io`, suggesting a cloaked crypto/NFT fraud route rather than confirmed malware delivery.

Its success against many filters likely came from that split design: the email itself looked like a plain transactional notice, the sender passed domain-level authentication, there was no attachment, and the suspicious destination was concealed behind a fingerprinting stage that static crawlers and simplified URL sandboxes could easily miss.

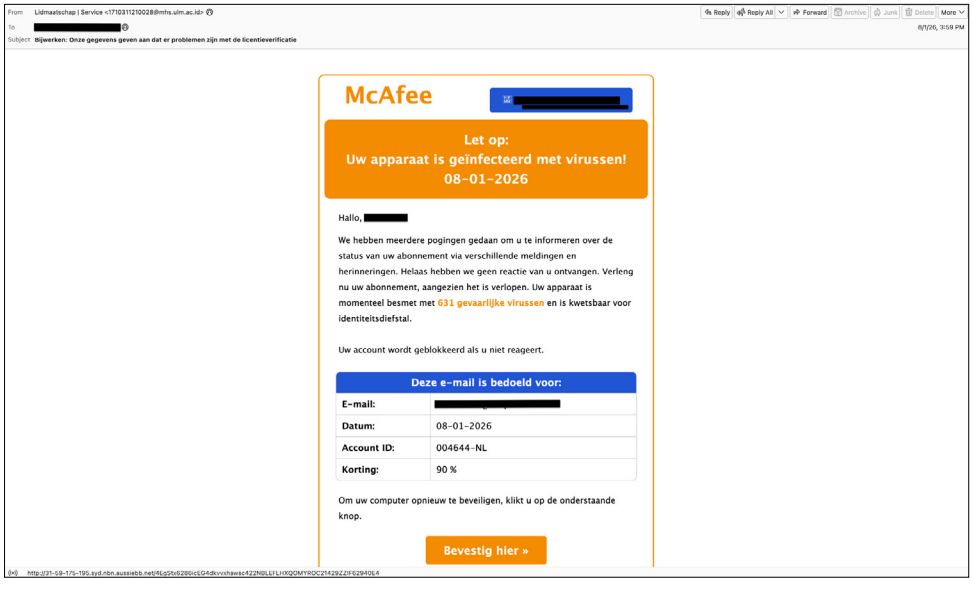

Antivirus renewal phishing sample.

Romanian PSD2 banking phishing hides behind IPv6-mapped URL obfuscation

A Romanian-language banking phishing sample observed in August 2026 impersonated *BCR S.A.* with a PSD2 'mandatory consent renewal' notice, warning recipients that access to web and mobile banking would be restricted unless they completed an online update by 27 August.

The message was sent from the DKIM-aligned but unrelated domain `xmasbrick[.]com` and contained no attachment, relying instead on a single call to action disguised as 'REÎNNOIRE CONSIMȚĂMÂNT' (consent renewal). The embedded link used IPv6-mapped URL notation, `hxxp://[0000:0000:0000:0000:0000:FFFF:67C1:B3DF]/11881659`, representing `103[.]193[.]179[.]223`, before redirecting through `web5-4s4c-online-garantibbva[.]vibtee[.]com/ro/` and eventually to *Google* during live verification.

The assessed threat is online-banking credential theft, with the active phishing endpoint likely cloaked, expired, or selectively serving benign content. This design made the sample harder to block because it paired convincing local regulatory language and bank branding with authenticated delivery, no executable payload, legitimate-looking footer details, and an obfuscated IP-literal URL format that can evade simplistic extraction, detonation, and reputation-based filtering.

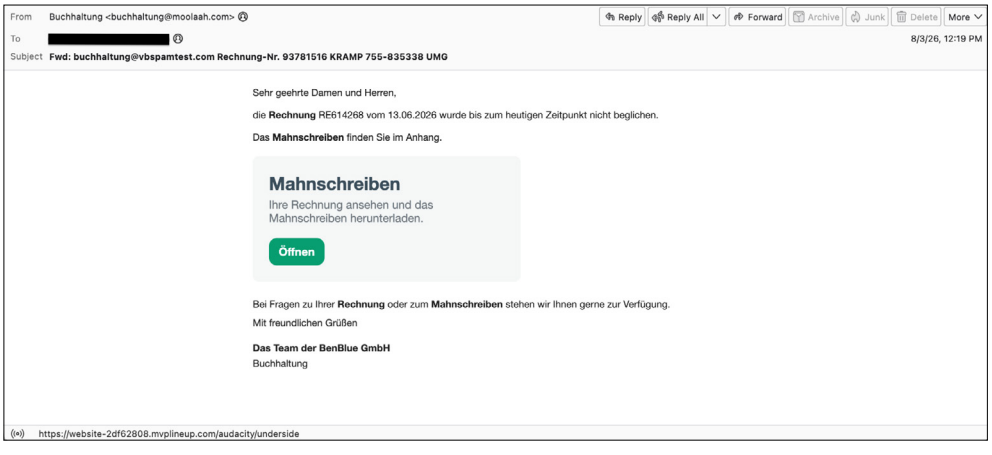

Invoice phishing sample redirecting to OpenSea.

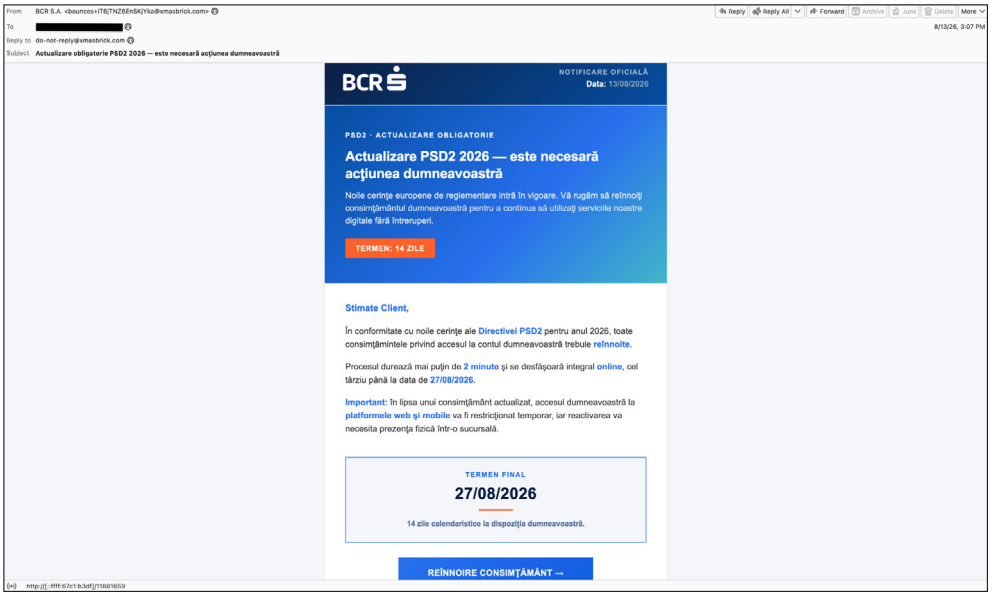

Romanian PSD2 banking phishing sample.

RESULTS

Of the participating full solutions, three achieved a VBSpam award – *Bluepex Mail Security, Coro Email Security* and *Zoho Mail* – while six others were awarded a 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: 99.994%
FP rate: 0.00%
Final score: 99.994
Malware catch rate: 100.000%
Phishing catch rate: 99.970%
Project Honey Pot SC rate: 99.995%
Abusix SC rate: 99.991%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Bitdefender delivered an outstanding performance, combining a spam catch rate of 99.994% with no false positives and a perfect malware catch rate. Phishing detection reached 99.970%, while results across all three spam feeds were exceptionally strong. Green speed measurements throughout secure the product another VBSpam+ award.

Bluepex Mail Security

SC rate: 99.887%
FP rate: 0.19%
Final score: 98.789
Malware catch rate: 99.970%
Phishing catch rate: 99.870%
Project Honey Pot SC rate: 99.864%
Abusix SC rate: 99.941%
MXMailData SC rate: 99.960%
Newsletters FP rate: 1.9%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Excellent detection was maintained by *Bluepex*, with 99.887% of spam, 99.970% of malware and 99.870% of phishing emails caught. However, false positives among legitimate mail and newsletters reduced the product's final score to 98.789, preventing it from retaining the VBSpam+ certification it achieved in the previous test. Nevertheless, consistently fast delivery and strong detection earn it a VBSpam award.

Coro Email Security

SC rate: 99.917%
FP rate: 0.37%
Final score: 98.088
Malware catch rate: 99.970%
Phishing catch rate: 99.910%
Project Honey Pot SC rate: 99.931%
Abusix SC rate: 99.871%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Coro Email Security produced a strong spam catch rate of 99.917%, including a perfect result on the *MXMailData* feed. Malware and phishing detection were also close to 100%, while delivery speeds remained green throughout. Although a 0.37% false positive rate brought the final score down, the product still earns VBSpam certification.

Fortinet FortiMail

SC rate: 99.960%
FP rate: 0.00%
Final score: 99.960
Malware catch rate: 100.000%
Phishing catch rate: 99.910%
Project Honey Pot SC rate: 99.963%
Abusix SC rate: 99.949%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Another excellent and well-balanced performance from *FortiMail* saw 99.960% of spam and every malware sample blocked, without a single false positive. The product also achieved consistently strong results across the three spam feeds, with all delivery-speed measurements remaining green. A final score of 99.960 earns the product another VBSpam+ award.

N-able Mail Assure

SC rate: 99.909%
FP rate: 0.00%
Final score: 99.909

Malware catch rate: 99.970%
Phishing catch rate: 99.930%
Project Honey Pot SC rate: 99.959%
Abusix SC rate: 99.770%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

N-able Mail Assure achieved a spam catch rate of 99.909% without blocking any legitimate emails or newsletters. Malware and phishing protection remained very strong, with every message from the *MXMailData* feed detected, although the *Abusix* stream proved slightly more challenging. Fast delivery speeds throughout complete a solid VBSpam+ performance.

N-able SpamExperts

SC rate: 99.910%
FP rate: 0.00%
Final score: 99.910
Malware catch rate: 99.970%
Phishing catch rate: 99.930%
Project Honey Pot SC rate: 99.959%
Abusix SC rate: 99.774%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Results for *N-able SpamExperts* closely matched those of its stablemate, with a 99.910% spam catch rate, no false positives, and a final score of 99.910. The product combined strong malware and phishing protection with perfect detection on the *MXMailData* feed and green delivery speeds throughout. The product comfortably earns another VBSpam+ award.

Net at Work NoSpamProxy

SC rate: 99.995%
FP rate: 0.00%
Final score: 99.995
Malware catch rate: 100.000%
Phishing catch rate: 99.990%
Project Honey Pot SC rate: 99.997%
Abusix SC rate: 99.989%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Net at Work NoSpamProxy recorded the strongest overall result in the test, combining a final score of 99.995 with no false positives. Every malware sample was blocked, 99.990% of phishing emails were caught, and perfect detection was achieved on the *MXMailData* feed. Exceptional consistency and green delivery speeds throughout complete the picture for a well-deserved VBSpam+ award.

Rspamd

SC rate: 57.873%
FP rate: 0.00%
Final score: 57.507
Malware catch rate: 61.440%
Phishing catch rate: 62.550%
Project Honey Pot SC rate: 58.896%
Abusix SC rate: 53.927%
MXMailData SC rate: 70.590%
Newsletters FP rate: 3.8%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Detection for *Rspamd* remained challenging, with only 57.873% of spam caught, although no false positives were recorded among legitimate ham emails. The product caught a little over 60% of both malware and phishing emails; the *Abusix* feed proved the most difficult, while *MXMailData* produced its strongest feed-level result. Delivery speeds were fast, but the final score of 57.507 remains well below the certification threshold.

Rspamd Premium 3.14.3

SC rate: 99.069%
FP rate: 0.75%
Final score: 95.410
Malware catch rate: 99.020%
Phishing catch rate: 98.730%
Project Honey Pot SC rate: 98.811%
Abusix SC rate: 99.871%
MXMailData SC rate: 97.830%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Rspamd Premium 3.14.3 showed a dramatic improvement over its performance in the previous test, raising its spam catch rate from 92.960% to 99.069% and making substantial gains in malware and phishing detection. However, a 0.75% false positive rate reduced the final score significantly. Despite much stronger detection and consistently fast delivery, the result remains short of the VBSpam certification threshold.

SEPPmail.cloudfilter

SC rate: 99.985%
FP rate: 0.00%
Final score: 99.985
Malware catch rate: 100.000%
Phishing catch rate: 99.970%
Project Honey Pot SC rate: 99.981%
Abusix SC rate: 99.996%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Another top-tier performance from *SEPPmail.cloudfilter* delivered a spam catch rate of 99.985%, no false positives, and a final score of 99.985. The product blocked every malware sample, achieved perfect detection on the *MXMailData* feed, and recorded the highest *Abusix* catch rate in the test, at 99.996%. Green delivery speeds throughout help secure another VBSpam+ award.

Zoho Mail

SC rate: 99.230%
FP rate: 0.00%
Final score: 99.047
Malware catch rate: 99.920%
Phishing catch rate: 99.590%
Project Honey Pot SC rate: 99.235%
Abusix SC rate: 99.149%
MXMailData SC rate: 99.930%
Newsletters FP rate: 1.9%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Zoho Mail achieved a respectable spam catch rate of 99.230% without producing any false positives among ordinary legitimate mail. Malware and phishing protection remained strong, *MXMailData* detection reached 99.930%, and the newsletter false positive rate improved from 3.8% to 1.9%. Although the later delivery measurements fell into the yellow bracket, the final score of 99.047 comfortably earns the product VBSpam certification.

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-ver30-1/.

The test ran for 16 days, from 12am on 1 August to 12am on 17 August 2026 (GMT).

The test corpus consisted of 103,969 emails. 103,380 of these were spam, 72,964 of which were provided by *Project Honey Pot*, 27,696 were provided by *Abusix* with the remaining 2,720 spam emails provided by *MXMailData*. There were 536 legitimate emails ('ham') and 53 newsletters, a category that includes various kinds of commercial and non-commercial opt-in mailings.

33 emails in the spam corpus were considered 'unwanted' and were included with a weight of 0.2; this explains the non-integer numbers in some of the tables.

Moreover, 3,563 emails from the spam corpus were found to contain a malicious attachment while 11,243 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[2].

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. 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

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

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
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	536	0	0.00%	6.6	103347	99.994%	99.994	SPAM + Verified
Bluepex Mail Security	535	1	0.19%	116.8	103236.8	99.887%	98.789	SPAM Verified
Coro Email Security	534	2	0.37%	85.6	103268	99.917%	98.088	SPAM Verified
Fortinet FortiMail	536	0	0.00%	41	103312.6	99.960%	99.960	SPAM + Verified
N-able Mail Assure	536	0	0.00%	93.6	103260	99.909%	99.909	SPAM + Verified
N-able SpamExperts	536	0	0.00%	92.6	103261	99.910%	99.910	SPAM + Verified
Net at Work NoSpamProxy	536	0	0.00%	5.4	103348.2	99.995%	99.995	SPAM + Verified
Rspamd	536	0	0.00%	43540.2	59813.4	57.873%	57.507	
Rspamd Premium 3.14.3	532	4	0.75%	962.2	102391.4	99.069%	95.410	
SEPPmail.cloudfilter	536	0	0.00%	15.2	103338.4	99.985%	99.985	SPAM + Verified
Zoho Mail	536	0	0.00%	795.4	102558.2	99.230%	99.047	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	0.0%	0	100.000%	3	99.970%	4	99.995%	2.6	99.991%	0	100.000%	0.09
Bluepex Mail Security	1	1.9%	1	99.970%	15	99.870%	99.4	99.864%	16.4	99.941%	1	99.960%	0.31
Coro Email Security	0	0.0%	1	99.970%	10	99.910%	50	99.931%	35.6	99.871%	0	100.000%	0.43
Fortinet FortiMail	0	0.0%	0	100.000%	10	99.910%	27	99.963%	14	99.949%	0	100.000%	0.16
N-able Mail Assure	0	0.0%	1	99.970%	8	99.930%	30	99.959%	63.6	99.770%	0	100.000%	0.31
N-able SpamExperts	0	0.0%	1	99.970%	8	99.930%	30	99.959%	62.6	99.774%	0	100.000%	0.31
Net at Work NoSpamProxy	0	0.0%	0	100.000%	1	99.990%	2.4	99.997%	3	99.989%	0	100.000%	0.09
Rspamd	2	3.8%	1374	61.440%	4210	62.550%	29983.6	58.896%	12756.6	53.927%	800	70.590%	15.39
Rspamd Premium 3.14.3	0	0.0%	35	99.020%	143	98.730%	867.6	98.811%	35.6	99.871%	59	97.830%	1.18
SEPPmail. cloudfilter	0	0.0%	0	100.000%	3	99.970%	14.2	99.981%	1	99.996%	0	100.000%	0.12
Zoho Mail	1	1.9%	3	99.920%	46	99.590%	557.8	99.235%	235.6	99.149%	2	99.930%	1.36

† *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	●	●	●	●
Bluepex Mail Security	●	●	●	●
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	
Net at Work NoSpamProxy	99.995
Bitdefender GravityZone Premium	99.994
SEPPmail.cloudfilter	99.985
Fortinet FortiMail	99.960
N-able SpamExperts	99.910
N-able Mail Assure	99.909
Zoho Mail	99.047
Bluepex Mail Security	98.789
Coro Email Security	98.088
Rspamd Premium 3.14.3	95.410
Rspamd	57.507

Hosted solutions	Anti-malware	IPv6	DKIM	SPF	DMARC	Multiple MX-records	Multiple locations
Bluepex Mail Security	ClamAV & Bitdefender	√	√	√	√	√	√
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					√