JUNE 2026

virus BULLETIN

Covering the global threat landscape

VBSPAM EMAIL SECURITY COMPARATIVE REVIEW JUNE 2026

Ionuţ Răileanu & Adrian Luca

In the Q2 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tested opted to be included in the public test, the rest opting for private testing (all details and results remaining unpublished).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.

This test set highlighted how email threats increasingly hide behind ordinary business workflows and reputable delivery infrastructure. The most notable samples we observed included a Danish domain-suspension lure sent through *Atlassian Jira*, a purchase-order attack that originated from a likely compromised *Microsoft 365* account and used a linked document, and an HR salary-adjustment message whose *Word* attachment led to a SilverFox payload chain. In each case, the malicious intent was pushed beyond the initial email, relying on authenticated senders, familiar notification formats and external payload staging to make gateway-level detection harder.

Despite this, most products maintained very high levels of protection across spam, malware and phishing, with several achieving near-perfect results and no false positives – though a smaller number of lower-scoring products did continue to struggle with these multi-stage, infrastructure-abuse techniques.

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 8,762 spam samples (8.96% 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	41.26%
2	China	14.84%
3	Japan	5.37%
4	Russian Federation	3.76%
5	Germany	2.27%
6	United Kingdom	2.07%
7	Netherlands	1.01%
8	Czech Republic	0.95%
9	Asia/Pacific region	0.93%
10	Ukraine	0.83%

Top 10 countries from which spam was sent.

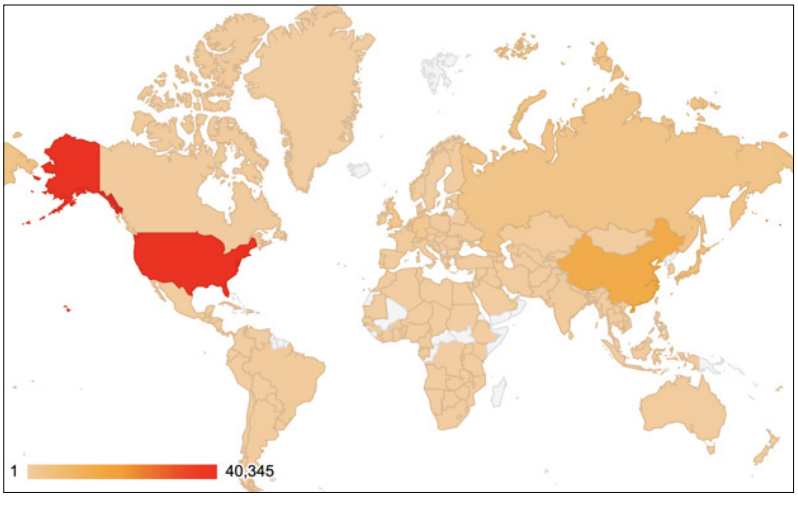

Geographical distribution of spam based on sender IP address.

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-TP196
- **Link**: https://www.amtso.org/tests/virus-bulletin-vbspam-q2-2026/

HIGHLIGHTS

Domain suspension phishing via Jira ticket notification

This message is a high-risk phishing email delivered through a legitimate *Atlassian Jira Service Management* notification flow. The attacker impersonates *Punktum dk A/S* and uses a Danish domain-suspension lure to pressure the recipient into 'reactivating' their `.dk` domain through a malicious `systeme[.]io` landing page.

The message is particularly challenging for security solutions because it inherits trust from *Atlassian* and *Amazon SES* infrastructure, passes DKIM for `atlassian[.]net`, and contains normal *Jira* notification headers, while the actual phishing content is embedded inside the service-desk message body.

This combination of valid authentication, trusted SaaS delivery and a low-reputation external CTA enables the message to bypass filters that rely heavily on sender reputation and platform-level trust.

Purchase order phishing via compromised Microsoft 365 account

This email is a high-risk phishing message sent from what appears to be a legitimate *Microsoft 365*-hosted business account, `Porcher Adams <porcher@wayupfront.co.za>`. The message uses a simple purchase-order lure

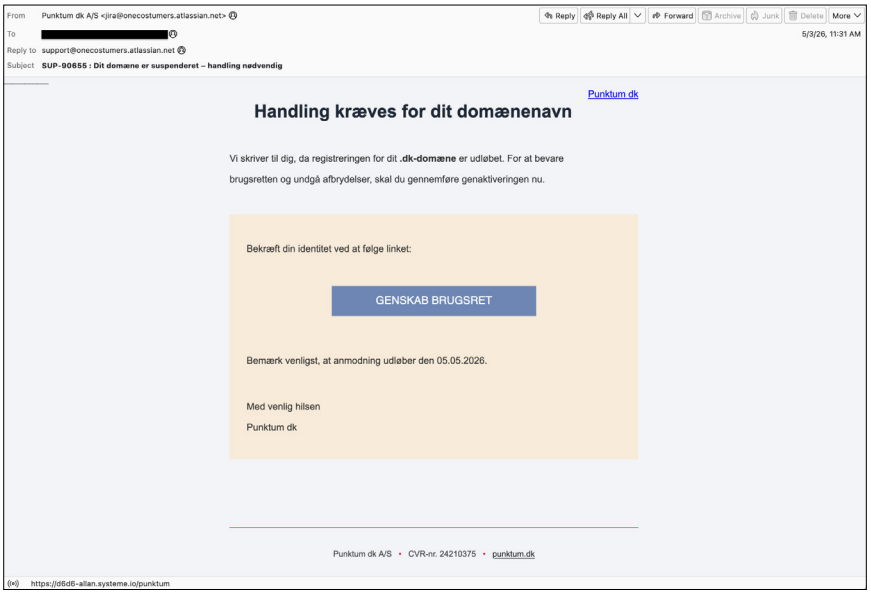

Phishing email impersonating Jira ticket notification.

Phishing sent via compromised M365 account

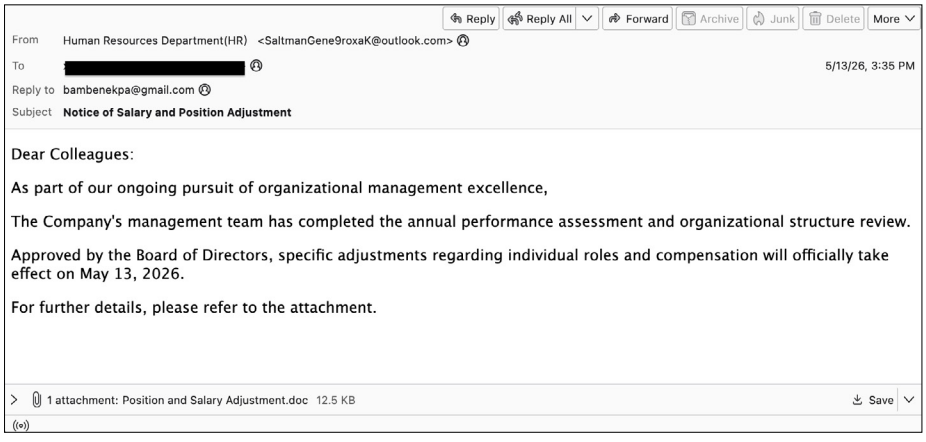

Malware sample luring to delivering ValleyRAT.

with the subject `Emailing: Customer PO` and asks the recipient to process a purchase order, attaching a *Word* document named `Purchase_Order.docx`.

The attachment is not macro-enabled, but it contains an external hyperlink labelled `VIEW PURCHASE ORDER HERE`, pointing to `hxxps://organisationdesign[.]github[.]io/labtechgroup/`, which is unrelated to the sender and serves as the next-stage phishing page[2].

This threat is challenging for security solutions because the message carries strong legitimacy signals: it was sent through *Microsoft 365* infrastructure, SPF/DKIM/DMARC passed for `wayupfront.co.za`, *Microsoft* classified it with low spam confidence (`SCL:1, SFV:NSPM`), and the email body itself is short, plain and business-like. The malicious element is pushed into an attached DOCX and then into an external link within that document, reducing the chance that content filters focused on the email body alone will detect the threat.

HR salary adjustment lure delivering ValleyRAT

This email is a high-risk malware delivery attempt using an HR-themed salary and position adjustment notice to push the recipient toward an attached *Word* document. The message is sent from a consumer *Outlook* account and uses a generic human resources display name, while the attachment contains a link to an external payload site rather than having malicious code directly embedded.

The attached *Word* document links to an external landing page that downloads a ZIP archive, `47cfcea4be4f89d9a25fabd7cd30e5469d7de839d6bded1a7f6760eec913a60a`, which after extraction and execution leads to a ValleyRAT[3] infection and contacts `akyv188[.]club`.

This makes the attack challenging for security solutions to detect, as the email itself is sparse, authenticated through *Microsoft* infrastructure, and the malicious behaviour is only exposed after the user opens the attachment and follows the embedded link.

RESULTS

Of the participating full solutions, one (*Zoho Mail*) achieved a VBSpam award, while eight others were awarded a VBSpam+ certification: *Bitdefender GravityZone Premium, Bluepex Mail Security, Coro Email Security, 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: 99.982%
FP rate: 0.00%
Final score: 99.982
Malware catch rate: 100.000%
Phishing catch rate: 99.940%
Project Honey Pot SC rate: 99.991%
Abusix SC rate: 99.920%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Bitdefender GravityZone Premium delivered an excellent performance in this test, combining a spam catch rate of

[2] https://app.any.run/tasks/8761b1f3-4bbb-4654-a9c9-5d13c43ff904
[3] https://www.proofpoint.com/us/blog/threat-insight/ta4922-suspected-chinese-crime-group-going-global

99.982% with no false positives and a final score of 99.982. It achieved a perfect malware catch rate, a phishing catch rate of 99.940%, and turned in strong results across all three spam feeds, while maintaining consistently fast delivery speeds throughout the test.

Bluepex Mail Security

SC rate: 99.898%
FP rate: 0.00%
Final score: 99.898
Malware catch rate: 99.940%
Phishing catch rate: 99.870%
Project Honey Pot SC rate: 99.886%
Abusix SC rate: 99.931%
MXMailData SC rate: 99.960%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Bluepex Mail Security produced a strong showing, with a spam catch rate of 99.898%, no false positives and a final score of 99.898. Its malware and phishing detection rates were both high, results across the *Project Honey Pot*, *Abusix* and *MXMailData* feeds were consistently solid, and all speed measurements remained comfortably in the green.

Coro Email Security

SC rate: 99.925%
FP rate: 0.00%
Final score: 99.925
Malware catch rate: 99.970%
Phishing catch rate: 99.840%
Project Honey Pot SC rate: 99.952%
Abusix SC rate: 99.730%
MXMailData SC rate: 99.990%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Coro Email Security performed well overall, achieving a spam catch rate of 99.925% with no false positives and a final score of 99.925. Malware and phishing catch rates were both high, though the *Abusix* feed proved slightly more challenging than the other streams, and while delivery speeds started strongly, later measurements moved into the yellow bracket.

Fortinet FortiMail

SC rate: 99.950%
FP rate: 0.00%
Final score: 99.950
Malware catch rate: 100.000%
Phishing catch rate: 99.950%
Project Honey Pot SC rate: 99.945%
Abusix SC rate: 99.954%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Fortinet's *FortiMail* turned in an excellent performance, with a spam catch rate of 99.950%, no false positives and a final score of 99.950. With perfect malware detection, a phishing catch rate of 99.950%, very strong feed-by-feed results and green speed measurements throughout, it combined accuracy with consistently prompt delivery.

N-able Mail Assure

SC rate: 99.974%
FP rate: 0.00%
Final score: 99.974
Malware catch rate: 99.970%
Phishing catch rate: 99.950%
Project Honey Pot SC rate: 99.997%
Abusix SC rate: 99.830%
MXMailData SC rate: 99.990%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

N-able Mail Assure achieved excellent results, combining a spam catch rate of 99.974% with zero false positives and a final score of 99.974. Malware and phishing detection were both very strong, the *Project Honey Pot* feed was handled particularly well. Its strong performance was backed up by fast delivery speeds across the board.

N-able SpamExperts

SC rate: 99.974%
FP rate: 0.00%
Final score: 99.974
Malware catch rate: 99.970%
Phishing catch rate: 99.950%
Project Honey Pot SC rate: 99.997%
Abusix SC rate: 99.830%
MXMailData SC rate: 99.990%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

N-able SpamExperts matched its stablemate, also posting a spam catch rate of 99.974%, no false positives and a final

score of 99.974. Its malware and phishing catch rates were high, feed-level performance was consistently strong, and all speed measurements remained in the green, resulting in a well-balanced overall showing.

Net at Work NoSpamProxy

SC rate: 99.988%
FP rate: 0.00%
Final score: 99.988
Malware catch rate: 100.000%
Phishing catch rate: 99.990%
Project Honey Pot SC rate: 99.990%
Abusix SC rate: 99.968%
MXMailData SC rate: 100.000%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Net at Work NoSpamProxy was one of the strongest performers in the test, with a spam catch rate of 99.988%, no false positives and a final score of 99.988. It achieved perfect malware detection, an excellent phishing catch rate of 99.990%, strong results across all spam feeds and consistently fast delivery times.

Rspamd

SC rate: 68.013%
FP rate: 3.31%
Final score: 51.594
Malware catch rate: 62.180%
Phishing catch rate: 54.810%
Project Honey Pot SC rate: 70.911%
Abusix SC rate: 62.853%
MXMailData SC rate: 47.540%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Rspamd had a difficult test, with a spam catch rate of 68.013%, a false positive rate of 3.31%, and a final score of 51.594. Malware and phishing detection both lagged well behind the leading products, with the *MXMailData* feed proving particularly challenging, and although delivery speeds were consistently fast, accuracy levels were not competitive.

Rspamd Premium 3.14.3

SC rate: 92.960%
FP rate: 0.17%
Final score: 92.118
Malware catch rate: 91.560%
Phishing catch rate: 94.140%
Project Honey Pot SC rate: 92.351%
Abusix SC rate: 95.459%
MXMailData SC rate: 94.860%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Rspamd Premium 3.14.3 improved substantially on the open-source version, reaching a spam catch rate of 92.960% and a final score of 92.118, though the false positive rate of 0.17% affected the overall result. Phishing detection slightly outperformed malware detection, the *Abusix* feed was handled especially well, and speeds remained solid throughout.

SEPPmail.cloudfilter

SC rate: 99.982%
FP rate: 0.00%
Final score: 99.982
Malware catch rate: 99.970%
Phishing catch rate: 99.970%
Project Honey Pot SC rate: 99.982%
Abusix SC rate: 99.992%
MXMailData SC rate: 99.970%
Newsletters FP rate: 0.0%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

SEPPmail.cloudfilter delivered excellent results, with a spam catch rate of 99.982%, no false positives and a final score of 99.982. Feed-by-feed performance was uniformly strong, phishing and malware detection were both close to perfect, and although the final speed measurement slipped into the yellow bracket, delivery was otherwise consistently fast.

Zoho Mail

SC rate: 99.385%
FP rate: 0.00%
Final score: 99.217
Malware catch rate: 99.970%
Phishing catch rate: 99.520%
Project Honey Pot SC rate: 99.369%
Abusix SC rate: 99.210%
MXMailData SC rate: 99.840%
Newsletters FP rate: 3.8%
Speed: 10%: ●; 50%: ●; 95%: ●; 98%: ●

Zoho Mail achieved a respectable spam catch rate of 99.385% and a final score of 99.217, with strong malware detection and no false positives on legitimate mail. However, a newsletter false positive rate of 3.8% and

slightly weaker performance on the *Abusix* and phishing streams weighed on the overall result, while later speed measurements also fell into the yellow.

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 2 May to 12am on 18 May 2026 (GMT).

The test corpus consisted of 99,011 emails. 97,781 of these were spam, 77,157 of which were provided by *Project Honey Pot*, 12,972 were provided by *Abusix*, and the remaining 7,651 spam emails were provided by *MXMailData*. There were 1,177 legitimate emails ('ham') and 53 newsletters (a category that includes various kinds of commercial and non-commercial opt-in mailings).

28 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, 3,461 emails from the spam corpus were found to contain a malicious attachment while 11,004 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[4].

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 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.

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

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	1177	0	0.00%	17.6	97741	99.982%	99.982	SPAM + Verified
Bluepex Mail Security	1177	0	0.00%	100.2	97658.4	99.898%	99.898	SPAM + Verified
Coro Email Security	1177	0	0.00%	73.4	97685.2	99.925%	99.925	SPAM + Verified
Fortinet FortiMail	1177	0	0.00%	48.4	97710.2	99.950%	99.950	SPAM + Verified
N-able Mail Assure	1177	0	0.00%	25	97733.6	99.974%	99.974	SPAM + Verified
N-able SpamExperts	1177	0	0.00%	25	97733.6	99.974%	99.974	SPAM + Verified
Net at Work NoSpamProxy	1177	0	0.00%	12	97746.6	99.988%	99.988	SPAM + Verified
Rspamd	1138	39	3.31%	31269.8	66488.8	68.013%	51.594	
Rspamd Premium 3.14.3	1175	2	0.17%	6882.6	90876	92.960%	92.118	
SEPPmail.cloudfilter	1177	0	0.00%	17.2	97741.4	99.982%	99.982	SPAM + Verified
Zoho Mail	1177	0	0.00%	600.8	97157.8	99.385%	99.217	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%	7	99.940%	7.2	99.991%	10.4	99.920%	0	100.000%	0.15
Bluepex Mail Security	0	0.0%	2	99.940%	14	99.870%	88.2	99.886%	9	99.931%	3	99.960%	0.24
Coro Email Security	0	0.0%	1	99.970%	18	99.840%	37.4	99.952%	35	99.730%	1	99.990%	0.26
Fortinet FortiMail	0	0.0%	0	100.000%	5	99.950%	42.4	99.945%	6	99.954%	0	100.000%	0.16
N-able Mail Assure	0	0.0%	1	99.970%	5	99.950%	2	99.997%	22	99.830%	1	99.990%	0.14
N-able SpamExperts	0	0.0%	1	99.970%	5	99.950%	2	99.997%	22	99.830%	1	99.990%	0.14
Net at Work NoSpamProxy	0	0.0%	0	100.000%	1	99.990%	7.8	99.990%	4.2	99.968%	0	100.000%	0.13
Rspamd	0	0.0%	1309	62.180%	4973	54.810%	22441	70.911%	4814.8	62.853%	4014	47.540%	10.77
Rspamd Premium 3.14.3	0	0.0%	292	91.560%	645	94.140%	5901	92.351%	588.6	95.459%	393	94.860%	5.81
SEPPmail. cloudfilter	0	0.0%	1	99.970%	3	99.970%	14.2	99.982%	1	99.992%	2	99.970%	0.11
Zoho Mail	2	3.8%	1	99.970%	53	99.520%	486.4	99.369%	102.4	99.210%	12	99.840%	0.69

† *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	99.988
SEPPmail.cloudfilter	99.982
Bitdefender GravityZone Premium	99.982
N-able Mail Assure	99.974
N-able SpamExperts	99.974
Fortinet FortiMail	99.950
Coro Email Security	99.925
Bluepex Mail Security	99.898
Zoho Mail	99.217
Rspamd Premium 3.14.3	92.118
Rspamd	51.594

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					√