AegisAI

BUYER'S GUIDE · 2026 EDITION

The Buyer's Guide to Email Security

How to choose email security when attackers write with AI. What to look for, who to involve and the questions to ask every vendor.

MICROSOFT 365 | GOOGLE WORKSPACE | VENDOR CHECKLIST INSIDE

CONTENTS

Inside this guide

Use it to plan a search, run a fair evaluation and compare vendors on evidence.

INTRO	**Introduction**	03
PART 1	**When to make a change**	04
PART 2	**Planning your search**	05
PART 3	**What to look for in an email security provider**	07
	01 · AI that investigates	07
	02 · Ease of use	09
	03 · Pricing and fees	10
	04 · Configurability and flexibility	11
	05 · Integrations	12
	06 · Layered security	13
	07 · Compliance and reporting	14
	08 · Scalability and performance	15
	09 · Reputation and support	16
END	**Conclusion**	17
PART 4	**Vendor checklist**	18
ABOUT	**About AegisAI**	19

Sources are named next to every figure. AegisAI figures are vendor-reported unless a customer or third party is named.

INTRODUCTION

The attack your team has to stop was written for one person.

Email is the one channel every employee has to leave open. For years, defenders had one advantage: attackers reused their campaigns. Someone saw an attack once, a rule was written, and the attack was blocked everywhere else.

Generative AI removed that advantage. A model can research a target from public sources, write a flawless message from a vendor they already work with, about a project that is real, and hide the payload behind a CAPTCHA. Doing it again for the next person costs almost nothing. Every message can be new.

Most email security in use today decides by recognition. Rules and reputation ask whether a message matches an attack seen before. Behavioral models ask whether it looks different from normal. An attack written for one person, from a trusted sender, in a normal tone, gives both of them nothing to find.

That changes what you should ask of a vendor. This guide covers when to make a change, how to plan your search, nine areas to evaluate and a checklist for comparing vendors side by side. Each area ends with questions to put to every vendor, AegisAI included.

It is written for security leaders, email and IT administrators, and the procurement and legal teams who work with them, at organizations running Microsoft 365 or Google Workspace.

AI SPEAR PHISHING, 2025

5x

growth in one year, from 2.8% to 13.9% of the phishing AegisAI observed.

Source: AegisAI production telemetry, 20,000+ phishing emails, start and end of 2025. aegisai.com

Q3 2026

+49%

rise in confirmed phishing, relative to overall email volume.

Source: AegisAI Q3 2026 State of the AI Attack Report. 100M+ emails across 1M+ protected mailboxes, July 1 to September 2, 2026. aegisai.com/state-of-ai-attacks-q3-2026

CLICK-THROUGH

54%

of people clicked fully AI-written spear phishing, the same rate as human experts. Generic phishing: 12%.

Source: Heiding et al., 2024, 101 participants. arxiv.org/abs/2412.00586

PART 1

When to make a change

Buyers usually arrive one of two ways. Their current tool missed an attack, or they can see the next wave coming and doubt the tool is ready. These are the signs it is time to look.

Attacks are reaching inboxes

Users report phishing your tools passed, or an incident review finds the first email sat in an inbox for days. Look at what got through: a real vendor on a real thread, a clean email with no attachment, a link that leads through redirects to a login page. Each is built to pass checks for known-bad senders and content.

Verdicts nobody can explain

When the tool blocks or passes a message, can your analyst see why? If the answer is a score and a list of signals, every escalation means redoing the work by hand. If the fix for a miss is a ticket to the vendor, the gap stays open while you wait.

Good mail stuck in quarantine

Rules tuned tight catch invoices, password resets and data requests, and someone releases them by hand. Tuned loose, real attacks land. If your team spends its week moving that dial, the approach is the problem.

Your environment changed

A move to Microsoft 365 or Google Workspace, an acquisition, new suppliers or new regions change who your people email and who attackers can pretend to be.

BUSINESS EMAIL COMPROMISE

$3.05B

in losses reported to the FBI in 2025, across 24,768 complaints. Up from $2.77B in 2024.

Source: FBI IC3, 2025 Internet Crime Report. ic3.gov

Tools that do not share

Account takeover, outbound data loss and phishing simulation often sit in separate products, consoles and renewals. Each handoff is a place one tool sees an attack and the next one misses it.

Cost that no longer adds up

Add the gateway, sandbox add-on, data loss prevention, training, rule upkeep and hours spent on reported email. One wire fraud can cost more than years of email security.

Renewals, a new security leader, an audit finding or a cyber insurance requirement are natural moments to compare. If two or more of these signs sound familiar, Part 2 covers how to plan the search.

PART 2

Planning your search

Start with the attacks you actually face, write requirements you can test, and bring the right people in early.

Identify your security needs

Pull the last quarter of user-reported email, incidents and near misses. Sort them by type: credential phishing, invoice and payment fraud, executive impersonation, QR codes, OAuth consent lures, calendar invites, account takeover. Note which ones your tools caught, which users reported and which you found late. Then list your compliance obligations, mail platforms, tenants and domains.

Define requirements and objectives

Turn each need into a test a vendor must pass on your own mail. For example: catches invoice fraud sent from a real, compromised vendor account. Shows the page behind a CAPTCHA. Removes a convicted email from every mailbox, forwards included. Keeps false positives inside a tolerance you set. Decide what must be true on day one and what can follow.

Identify stakeholders

Email security touches several teams. Bring each in before you shortlist.

Security leadership	Owns the risk and the board conversation. Needs a clear difference from today and verdicts the team can defend.
Security operations	Lives with the alerts. Needs fewer escalations, evidence on every verdict and fast removal.
IT, messaging and architects	Own Microsoft 365 or Google Workspace. Review permissions, data flow and what happens when the service fails.
Procurement, legal and privacy	Review terms, data processing, security evidence and pricing.
Finance and executives	Weigh cost and exposure. They are also the people attackers target most for payment fraud.
User representatives	Department heads and frequent targets report how banners, released mail and reporting feel day to day.

P2 PLANNING YOUR SEARCH · CONTINUED

Plan a proof of value

Vendor claims are hard to tell apart: every vendor now says it uses AI, and most say they explain their verdicts. The fastest way through is to run the product on your own mail before you commit. A fair proof of value connects read-only, looks back over recent mail, runs beside your current tools and shows what each one caught. Agree in writing before you start: which attack types count, who labels the results, how false positives are measured, how long it runs and who signs off.

HOW AEGISAI RUNS IT · FIVE DAYS



DAY 0

Connect

An admin approves the API connection in about five minutes. No MX change.

DAYS 1 TO 4

Read-only scan

New mail and a look-back over recent mail. Nothing is moved or deleted.

DAY 5

Findings review

What AegisAI caught that your stack missed, each with its evidence.

WHEN READY

Turn on removal

Grant write access. You can revoke it at any time.

You keep the findings report either way. Source: aegisai.com homepage FAQ and AegisAI one-pager, 2026.

PLANNING CHECKLIST	DONE
Identify security needs	
Review last quarter's reported email, incidents and near misses	☐
Map the attack types that reached users, and how each was found	☐
Confirm compliance requirements, platforms, tenants and domains	☐
Define objectives	
Write requirements as tests a vendor must pass on your mail	☐
Prioritize: day-one must-haves and what can follow	☐
Agree proof-of-value success criteria in writing	☐
Identify stakeholders	
Security leadership and security operations	☐
IT, messaging administrators and security architects	☐
Procurement, legal, privacy and finance	☐
User representatives	☐

PART 3

What to look for in an email security provider

Evaluate vendors across nine areas. Each ends with questions to ask.

01 AI-ENABLED DETECTION

AI that investigates, not just scores

Every vendor says it uses AI. Ask what the AI does with an email: each approach asks a different question.

RULES AND REPUTATION

Does this match a known attack?

Misses a new campaign, a clean sender and a message with nothing to match.

BEHAVIORAL BASELINES

Is this different from normal?

Misses a hijacked account writing from inside its own history, and a first contact with no history at all.

INVESTIGATION

Is every claim in this email true?

Checks evidence the sender does not control: where links really lead, whether the two parties have written before, headers and sending infrastructure.

HOW AEGISAI INVESTIGATES AN EMAIL

01 · CONNECT

Reads every email by API. No MX change.

02 · INVESTIGATE

Specialist agents check sender and identity, links and QR codes, attachments, intent and urgency, OAuth lures and payment requests.

03 · DECIDE

Proteus, AegisAI's own model, weighs every claim against the evidence. Not settled? The agents keep digging.

04 · ACT

The verdict ships with its reason. Convicted mail is removed from every mailbox it reached.

MALICIOUS **Vendor payment fraud**

- ✓ Lookalike sender domain, registered days ago
- ✓ Reply-to address does not match the vendor
- ✓ First bank-detail change in the thread's history

Removed from every inbox. *Illustrative example. Names and details are fictional.*

01 AI-ENABLED DETECTION · CONTINUED

Go through the gate

Attackers hide payloads behind CAPTCHAs, bot checks, redirect chains and links inside PDFs, because a sandbox that cannot pass the gate sees a blank page and calls it clean. Ask whether the vendor follows each link hop by hop to the final page. AegisAI opens links, files and QR codes in a sandbox and follows them to the end. A page it cannot reach is treated as evidence, never scored clean.

Evidence you can read

A verdict should arrive with the steps taken and what each found, so an analyst reviews the investigation instead of redoing it. Every AegisAI verdict carries its reason in plain language.

Models built for email

General-purpose models follow the instructions in front of them, which is what phishing exploits, and attackers can test public models privately until a message slips through. AegisAI trains and runs its own models. Each new version of Proteus must pass adversarial tests across nine attack families before release.

False positives count too

Blocking good mail has a cost: vendors, customers and executives waiting for someone to release their messages. Ask how each vendor measures false positives in production, not only catch rates.

IN PRODUCTION **99.996%**

observed precision for Proteus verdicts.

Source: AegisAI, Proteus launch post, October 2026.

PROTEUS BENCHMARK

>99%

of 3,020 held-out phishing attacks identified. A leading closed-source general model: 81.3%. A widely used open-weight model: 80.0%.

7.2x faster median verdict than the closed-source model. Human-reviewed attacks from one week of AegisAI production traffic. Counts missed detections only; the general models received the message and a plain classification instruction. Source: AegisAI, Proteus launch, October 8, 2026

ASK EVERY VENDOR

- ? Take the last phishing email you blocked. Can you show the page at the end, behind the CAPTCHA?
- ? Can you show why, step by step, without a support ticket?
- ? Whose model makes the verdict, and what was it built to do?
- ? When a new attack type appears, who retrains the model, and how long does it take?
- ? What decides that an investigation is finished?

02 EASE OF USE

Quiet for users. Light for administrators.

Good email security disappears into the working day. Test it from three seats: the end user, the administrator and the analyst.

Nothing to maintain

Ask how many rules, allow lists, block lists and keyword filters the product needs to stay effective, and who writes them. AegisAI has no rule library to maintain and no baseline period to wait out. Mesh reports zero manual tuning since deploying it.

A clean inbox without a quarantine queue

Newsletters, promotions and spam should move aside as they arrive, with quarantine kept for real threats. Bulk mail deserves the same scrutiny as everything else, because phishing hides in it.

Clear signals on flagged mail

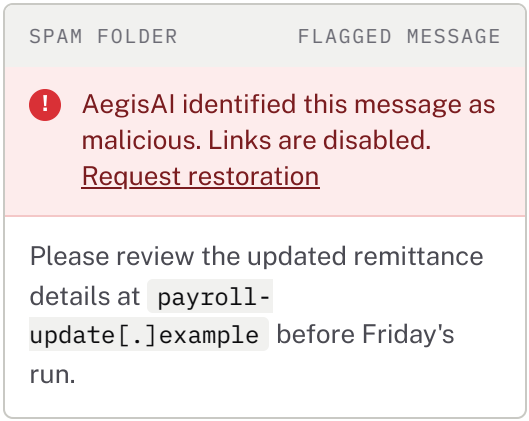
If a flagged message stays reachable, the user should see why and should not be one click from a live link. In AegisAI's banner modes, the message moves to spam with a warning banner, links shown as plain text, tracking pixels removed and attachments stripped. The original is kept, so an administrator can restore it intact.

Reporting that closes the loop

A reported email should get the same investigation as everything else, and the person who reported it should hear back. AegisAI investigates every reported message and replies to the reporter with the verdict.

Analysts confirm instead of investigating

Each alert should arrive with a summary, the evidence and a next step. The AegisAI console renders every message in an isolated sandbox, with its links, decoded QR codes, redirect chain and final page.



Illustrative example. Names and details are fictional.

> "Ryan set it up within five minutes. That's it. It just works."
>
> ARTHUR STROMQUIST, SECURITY LEAD, LANGCHAIN

ASK EVERY VENDOR

- ? How many rules or lists will we maintain, and who writes them?
- ? What does a user see on a flagged message, and how is one released?
- ? What happens after a user reports an email? Does the reporter hear back?
- ? What does an analyst see on a typical alert?

03 PRICING AND FEES

Price the stack, not the license.

The cheapest license can sit on top of the most expensive stack. Price what you will actually run, and the time your team spends running it.

Know the licensing unit

Most email security is priced per mailbox or per user, per year. Ask what counts: shared mailboxes, inactive users, aliases and service accounts. Ask what happens when someone leaves, what you pay if you grow past the licensed count and whether you can step down at renewal.

Count every line item

List what you pay for separately today: the gateway, sandboxing, data loss prevention, awareness training, account takeover protection and any managed triage service. Then add the hours spent writing rules, releasing quarantined mail and triaging reported email, plus soft costs: implementation services, premium support and extra consoles.

Ask how AI costs reach you

If a vendor builds on rented general-purpose models, ask how a change in token prices reaches your bill. A vendor that owns its models controls that cost.

Prove value before you pay

A proof of value on your own mail shows what you would be paying for. AegisAI's runs in five days, and you keep the findings report either way.

HOW AEGISAI PRICES

- ✓ **Per protected mailbox, per year, by module.**
- ✓ Inbound email security covers phishing, malware and business email compromise, spam and graymail filtering, and triage of user-reported email.
- ✓ Account takeover, outbound DLP, shadow IT discovery, and simulation and training add on the same connection, in the same console, under one contract.
- ✓ Standard support is included. Usage above the licensed count is billed pro rata at your contracted rate.

ASK EVERY VENDOR

- ? What is the licensing unit, and what counts toward it?
- ? Which capabilities are included, and which are add-ons?
- ? What happens if we exceed or fall below our count?
- ? Are there deployment, onboarding or support fees?

04 CONFIGURABILITY AND FLEXIBILITY

Your team decides what happens next.

Control over the action matters as much as the verdict: how you deploy, how you roll out and what happens to what it finds.

Deployment

A gateway needs an MX change, and every email routes through it. An API-based product reads mail after delivery, outside the delivery path. AegisAI is API-based and cloud-only: no MX change, and it runs beside a gateway or native filtering.

Roll out at your own pace

Start read-only, measure, then act. AegisAI starts in monitoring mode, with no time limit. Removal needs write access your admin grants and can revoke. Most customers move up the ladder below, and can move back down.

AEGISAI MODE	WHERE THE FLAGGED MESSAGE GOES	WHAT THE USER SEES	WHO RESTORES IT
1 **Read-only**	Stays where it was delivered	Nothing changes	N/A
2 **Soft quarantine**	Spam folder	The message, unchanged	User
3 **Banner and defang**	Spam folder	Warning banner, links disabled	User
4 **Admin restoration**	Spam folder	Banner, links disabled, a link to request restoration	Administrator
5 **Full quarantine**	Out of the mailbox	A notice in the spam folder	Administrator

Spam folder means Junk Email in Microsoft 365 and Spam in Gmail.

Scope and exceptions

Exclude specific mailboxes, or add trusted domains to a custom skip list so their links are not crawled. Every other check still runs.

Policies where you need them

Outbound DLP, in early access, starts from compliance packs or detector rules scoped by sender and recipient, with an audit mode for live mail.

ASK EVERY VENDOR

- ? Does it sit in our mail flow? What changes in routing?
- ? Can we run read-only first, for as long as we need?
- ? Which actions can it take on a verdict, and who can restore a message?
- ? What still runs on excluded mailboxes or trusted domains?

05 INTEGRATIONS

Fit the stack you already run.

Email security should connect where your mail and identities live, send what it finds to your security tools and act where the account is.

MAIL PLATFORMS

Microsoft 365 through Microsoft Graph, with a publisher-verified Entra application for each customer and no stored client secret. **Google Workspace** through the Gmail API with domain-wide delegation. Read-only permissions for the evaluation; one added permission turns on quarantine.

IDENTITY

Okta and **Microsoft Entra ID** supply the directory, sign-in signals for account takeover, targeting for simulations and approved apps for shadow IT. Console sign-in runs through Google Workspace or Microsoft Entra ID over OpenID Connect, so AegisAI never stores user passwords.

SIEM

Verdicts and indicators export to your SIEM. A documented connector for **CrowdStrike Falcon Next-Gen SIEM** sends each detection with its explanation and indicators, with first events within 10 to 15 minutes of setup.

ALERTS AND CHAT

A **Slack** app that an administrator connects from the AegisAI console. Choose which alerts reach which channel.

USER REPORTING

The report button in **Outlook** and **Gmail**, the AegisAI Report Phishing add-on for Google Workspace, or your abuse mailbox. Reports get the full investigation.

API AND YOUR AGENTS

A read-only **REST API**, enabled on request, with scoped, revocable tokens to pull detections, verdicts and their reasoning, and an **MCP server** so your own agents can query the same data.

Training history

Moving simulations to AegisAI? Import past campaign history from your current platform so your trend lines continue.

ASK EVERY VENDOR

- ? Which permissions do you need for the evaluation, and which for enforcement?
- ? Which SIEM, SOAR and ticketing tools do you connect to today, not on the roadmap?
- ? Is there an API, and what can it read and write?

06 LAYERED SECURITY

Layers that pass evidence to each other.

No single check stops every attack. The layers that matter share what they find, from the first email to a compromised account.

L1	**Native filtering** Blocks bulk spam, known malware and known-bad links before delivery.	**YOUR PLATFORM** MICROSOFT OR GOOGLE
L2	**Investigation of every delivered email** Agents check sender, intent, links, attachments and payment requests. Convicted mail is removed from every mailbox it reached, forwards and calendar invites included.	**AEGISAI** AVAILABLE
L3	**User reports, investigated** Every reported email gets the same investigation, and the reporter hears back.	**AEGISAI** AVAILABLE
L4	**Account takeover monitoring** Sign-ins and mailbox activity weighed together. Findings come with recommended steps; your team acts in its identity provider.	**AEGISAI** EARLY ACCESS
L5	**Outbound data loss prevention** Personal data, financial data and secrets in outbound mail, judged by who is sending and to whom.	**AEGISAI** EARLY ACCESS
L6	**People: simulation and training** Real attacks caught in your tenant, replayed as training to the people they targeted. An AI red team agent tests exposure with your approval.	**AEGISAI** AVAILABLE

WHY LAYER 2 EXISTS

3 in 10

confirmed phishing emails cleared native Gmail and Microsoft filters in Q3 2026.

Source: AegisAI Q3 2026 State of the AI Attack Report.

Strong encryption

Data encrypted in transit with TLS 1.2 or later and at rest with AES-256. AegisAI also supports customer-managed encryption keys.

Threat intelligence that keeps up

Models retrained continuously. A new version replaces the current one only after adversarial testing and a side-by-side run on live traffic.

Ask every vendor which layers they cover, which are generally available today and whether the layers share findings.

07 COMPLIANCE AND REPORTING

Hold your vendor to your own standard.

Your email security vendor processes every message your company sends and receives. Ask for evidence, not adjectives.

Attestations and testing

Ask for the SOC 2 report and its audit period, any ISO certification and its scope, and the date of the last independent penetration test. AegisAI holds SOC 2 Type II and is working toward ISO 27001. Independent penetration tests run at least once a year. AegisAI signs Business Associate Agreements with healthcare customers, and its data processing addendum includes EU Standard Contractual Clauses.

Data handling

Ask what is stored, where and for how long, and whether any of it trains a model. At AegisAI, mail judged safe is analyzed in memory and discarded. Content of flagged messages is scanned for sensitive data before storage, encrypted and kept separately for each customer, in the United States. Customer email is never sent to outside LLM APIs. AegisAI trains its own models only on content its system and analysts both confirm as malicious. Customer-managed keys and bring-your-own-cloud deployment are available.

Reporting and audit

Every AegisAI verdict keeps its evidence: what was checked, what was found and the final page. Threat history stays in the console for 90 days by default, with longer retention on request. DLP findings stay masked until an admin opens them, and every red team lure and event is exported as an audit record.

YOUR EMAIL SECURITY SHOULD ADDRESS

- ✓ Credential phishing
- ✓ Business email compromise and vendor fraud
- ✓ Malware and ransomware delivery
- ✓ Account takeover
- ✓ QR code, OAuth and calendar lures
- ✓ Spoofing, with SPF, DKIM and DMARC treated as one input, not a verdict
- ✓ Spam and graymail
- ✓ Outbound data loss
- ✓ Shadow IT and AI apps
- ✓ Simulation and training

ASK EVERY VENDOR

- ? Is our email used to train any model? Is that opt-in, and is it in the contract?
- ? Who are your subprocessors, and where is data processed?
- ? Will you sign our DPA, or a BAA?

08 SCALABILITY AND PERFORMANCE

Fast enough that users never meet the attack.

An API-based product reads mail after delivery, so speed is a security question: the gap between arrival and removal is the attacker's window.

1.9s

Median time from arrival to verdict, link detonation included. Removal follows within seconds.

Source: aegisai.com product pages, 2026.

Time to verdict and removal

Ask for the median and worst-case time from arrival to verdict, and from verdict to removal across every mailbox. A slow verdict on an API-based product is time a user can click.

Mail keeps flowing

A gateway in the mail path can delay or stop mail when it is slow or down, which matters for time-sensitive mail such as one-time passcodes. An API-based product sits outside that path. If AegisAI is unavailable, mail keeps flowing and detection pauses until it returns.

AEGISAI STANDARD SERVICE LEVELS

99.9%
monthly availability commitment

99%
of emails processed within 30 seconds

5 min
to remove a threat found after delivery

Recovery

Recovery targets are two hours to restore service and no more than 15 minutes of data loss. A July 2026 live restore test recovered in 21 minutes.

Volume

In Q3 2026, AegisAI analyzed more than 100 million emails across more than one million protected mailboxes.

ASK EVERY VENDOR

- ? What are your median and 99th-percentile times to verdict?
- ? What happens to our mail if your service goes down?
- ? What uptime does the SLA commit to, and what is the remedy?
- ? What are your recovery objectives, and when were they last tested?

09 REPUTATION AND SUPPORT

Check the team before you sign.

Support decides how the contract feels in year two. So does how fast the vendor's product moves.

Track record

Ask who built the product and what they protected before. AegisAI's CEO, Cy Khormaee, spent five years at Google leading the teams behind Safe Browsing, reCAPTCHA and Web Risk. CTO Ryan Luo spent nine years at Google and modernized the core Safe Browsing phishing platform. The research team includes former leads of Chrome Enterprise Security, Google Ads Security and GCP AI Platform Security. AegisAI has raised $49M, including a $36M Series A led by Battery Ventures in July 2026, with Accel and Foundation Capital.

Implementation

Ask how much of the work falls on your team. Connecting AegisAI takes an administrator about five minutes and changes nothing in mail flow. AegisAI runs the proof of value and the findings review.

Ongoing support

Ask for response times in writing. AegisAI's standard support, included with every subscription, commits to a first response within one hour, around the clock, when the service is down, and within four business hours when it is degraded.

Product velocity

Ask what the vendor shipped in the last six months, which of it you would pay for separately, and whether each module is generally available, in early access or on the roadmap.

> "Mesh is a fast-moving, tech-forward organization. We need vendors like [AegisAI] that innovate with us in our highly dynamic environment."
>
> DANIEL HOOPER, CISO, MESH

CUSTOMER PROOF · MESH

300+

phishing and impersonation emails the existing tools had missed, surfaced in the first scan.

Mesh also reports zero missed legitimate business emails and zero manual tuning since deploying AegisAI. Source: aegisai.com/customers.

ASSESS EVERY VENDOR'S

- ✓ Expertise and track record
- ✓ Customer references you can call
- ✓ Response times, in writing
- ✓ Product velocity and module status
- ✓ Independent attestations

CONCLUSION

Choose on evidence.

AI has made the targeted attack cheap. A message written for one person, from a sender they trust, with the payload behind a gate, is now routine. Email security built to recognize yesterday's attacks will keep missing it.

When you compare vendors, start with how each one decides: whether it investigates or matches, whether it goes through the gate, and whether its verdicts arrive with evidence. Then weigh ease of use, total cost, control over actions, integrations, layered coverage, compliance, speed and support.

Above all, test on your own mail. A read-only proof of value, run beside your current tools, shows what is getting through today.

See what your current stack is missing.

Connect a tenant in monitoring mode and compare in five days.

Book a demo →

FURTHER READING ON AEGISAI.COM

AegisAI Platform Overview
aegisai.com/platform

Q3 2026 State of the AI Attack Report
aegisai.com/state-of-ai-attacks-q3-2026

PART 4

Vendor checklist

Build your shortlist or make the final call. Ask each vendor to show you, not tell you.

CHECKLIST FOR SELECTING AN EMAIL SECURITY PROVIDER	AEGISAI	VENDOR 1	VENDOR 2
Protection			
Catches AI-written phishing, BEC and vendor fraud	Investigates every email	☐	☐
Follows links past redirects and CAPTCHAs	Sandbox follows every hop	☐	☐
Explains every verdict with evidence	Reason on every verdict	☐	☐
Models built for email, retrained by the vendor	Proteus, its own model	☐	☐
Account takeover, DLP and training in one console	Yes; ATO and DLP in early access	☐	☐
Performance and reliability			
Mail keeps flowing if the service is down	API, out of the mail path	☐	☐
Fast verdict and removal from every mailbox	1.9s median to verdict	☐	☐
Uptime commitment in the contract	99.9% availability SLA	☐	☐
Ease of use			
No rules or lists to write and maintain	No rule library	☐	☐
Reported email investigated, reporter answered	Yes	☐	☐
Read-only start and a choice of actions	Monitoring mode, five modes	☐	☐
Compliance			
Independent security attestation	SOC 2 Type II	☐	☐
Clear data handling; no outside LLM APIs	US processing, CMEK	☐	☐
Pricing and fees			
Clear licensing unit, inclusions and overage	Per mailbox, by module	☐	☐
Proof of value on your own mail	Five days; keep the report	☐	☐
Support			
Response times committed in writing	1 hour, 24/7, for outages	☐	☐
Credible team and customer references	Google Safe Browsing founders	☐	☐

AegisAI

ABOUT AEGISAI

The first AI lab for email security

Every email, investigated.

AegisAI puts an investigation on every email. Its agents check who really sent a message, follow every link to the final page and read what the message is asking someone to do. Proteus, AegisAI's own model, weighs every claim against the evidence, and every verdict ships with its reason.

AegisAI connects to Microsoft 365 and Google Workspace by API in about five minutes, with no MX change, and runs beside your current tools. Its founders led Google's Safe Browsing and reCAPTCHA teams, and it is backed by Battery Ventures, Accel and Foundation Capital.

5-MINUTE INSTALL | PROVEN IN 5 DAYS | NO MX CHANGE | SOC 2 TYPE II

Book a demo →

© 2026 AEGISAI SECURITY, INC.