

2026

seceon

Beyond the Filter: Why Email Security Is Becoming an AI-Driven Security Operations Discipline

How detection, correlation, and response are redefining the email security market and where Seceon aiSecurity Email360 fits

Executive Summary

Email remains the single most exploited entry point into the enterprise. It is also the channel security leaders understand least in real time most organizations can tell you whether a message was blocked or delivered, but far fewer can tell you which users are being targeted, how threats are trending week over week, or how an email-borne compromise connects to what happens next on the endpoint or in the cloud.

The email security market has tracked a clear trajectory over the last several years: native cloud email platform defenses have matured, attackers have adapted with more sophisticated, AI-assisted social engineering, and the category has correspondingly shifted from simple secure email gateways toward API-integrated, cloud-native platforms that layer behavioral analysis, threat intelligence, and post-delivery detection on top of native controls. Industry-wide, the consistent view is that no single control is sufficient effective programs combine platform-native protections with additional layers of detection and response.

This whitepaper examines that evolution: what is driving it, what capabilities security and risk leaders should expect from a modern email security investment, and how the discipline is converging with the broader security operations center (SOC) rather than remaining a standalone mailbox control. It closes by describing how Seceon aiSecurity Email360 applies these principles turning email telemetry into a structured, investigable stream of security signal, using a Connect → Collect → Analyze → Correlate → Act workflow to move organizations from passive filtering toward active visibility, detection, investigation, and response.

> ***Key takeaway:*** *Email security is no longer a filtering problem. It is a visibility, correlation, and response problem and buyers should evaluate solutions accordingly.*

The Email Security Blind Spot Hasn't Closed It Has Moved

For two decades, the email security conversation centered on stopping spam and known-bad attachments at the gateway. That problem is, for the most part, solved. What has replaced it is harder: attackers who don't need malware at all, who impersonate a trusted colleague or vendor, and who exploit the fact that a single well-timed, well-written message can move money, credentials, or sensitive data without ever tripping a signature-based filter.

The scale of the problem is now well documented by independent sources. The FBI's Internet Crime Complaint Center (IC3) reported that business email compromise (BEC) generated roughly $3.05 billion in confirmed U.S. losses in 2025, up from about $2.77 billion the year before making BEC the second-most costly category of cybercrime the Bureau tracks, behind only investment fraud. The average confirmed loss per BEC complaint now exceeds $120,000, reflecting a pattern in which attackers favor a small number of high-value targets over high-volume, low-value spam (FBI IC3, 2025 Internet Crime Report).

Verizon's Data Breach Investigations Report tells a parallel story from the breach-investigation side: the human element phishing, pretexting, credential misuse, and simple error continues to factor into roughly 60% of confirmed breaches, and email is consistently cited as one of the top initial-access vectors alongside stolen credentials and exploited vulnerabilities (Verizon, Data Breach Investigations Report, 2025). Put simply: most breaches still start with a message in an inbox.

What has changed most in the last two years is not the target but the tooling attackers use to reach it. Microsoft's 2025 Digital Defense Report found that AI-automated phishing campaigns achieved click-through rates of roughly 54%, compared with about 12% for traditional phishing a more than four-fold increase in effectiveness and noted that BEC now accounts for a larger share of successful attack outcomes than ransomware (21% versus 16%) even though it represents a much smaller share of overall attack volume (Microsoft Digital Defense Report, 2025). Generative AI has lowered the cost of writing a convincing, grammatically clean, contextually accurate lure, and it is increasingly being used to clone voices, mimic writing style, and localize social-engineering content for different regions and languages.

The practical implication for security leaders is that email risk can no longer be managed as a mailbox-hygiene problem owned solely by IT operations. It requires the same discipline telemetry, detection logic, investigation workflow, and measurable response that organizations already apply to endpoint and network security.

Defining the Modern Email Security Category

Email security has evolved from a niche, spam-focused product category into one of the most strategically important areas of enterprise security spend. That evolution has been driven by three converging forces: the near-total migration to cloud office platforms, the rise of conversational, payload-free attacks like BEC that evade traditional filtering, and the measurable improvement in attacker tooling brought on by generative AI. Understanding how the category itself has changed is the first step to evaluating any vendor within it.

The category has moved toward API-based, cloud-native architectures

As organizations continue migrating to cloud office platforms such as Microsoft 365 and Google Workspace, native platform defenses have improved substantially, and supplemental email security solutions increasingly connect via API rather than sitting as an inline mail gateway. This shift reduces deployment friction and mail-flow risk, but it also means buyers must evaluate how quickly a solution can act quarantine, notify, or block once it detects a threat, since API-based tools generally operate after a message has already reached the mailbox.

Detection quality is difficult to compare across vendors so layering is common practice

The sheer volume and sophistication of email-enabled social engineering attacks, combined with the practical difficulty of consistently measuring true detection efficacy across vendors, has led many organizations to run more than one layer of email defense rather than relying on a single control. For buyers, this reinforces two things: first, that any single vendor's detection claims should be treated as one input among several, not a final verdict; and second, that a platform's ability to integrate with and add value on top of existing controls (native cloud defenses, a secure email gateway, or a SIEM/XDR platform) is itself a differentiator, not a compromise.

The category is converging around behavioral and identity-centric detection

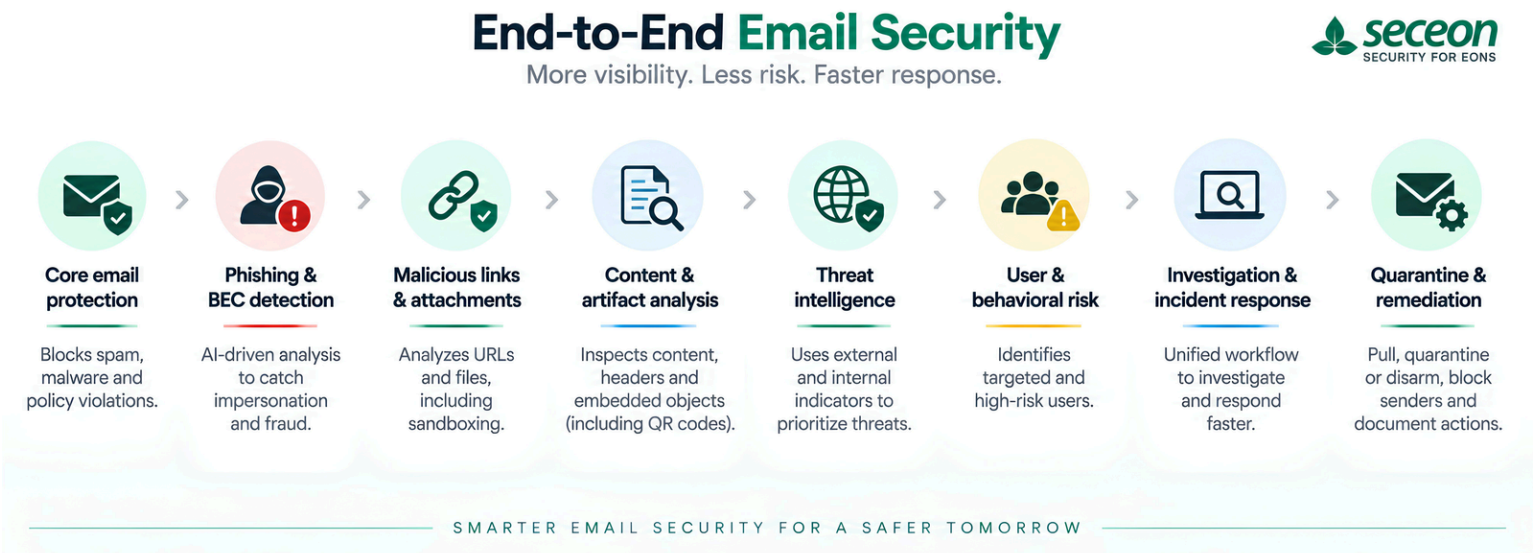
Because modern BEC and impersonation attacks frequently contain no malicious link or attachment at all, the category has shifted emphasis toward vendors' ability to model normal communication behavior sender relationships, writing style, financial-request patterns, and domain history rather than relying solely on signature or reputation-based detection. This is consistent with what independent breach and fraud data shows: the attacks causing the largest financial losses today are conversational, not malware-based (see Section 1).

Taken together, this market direction supports a simple filter for evaluating any email security investment: does it extend visibility beyond the inbox, does it detect behaviorally rather than only signature-based threats, and does it connect email findings to the rest of the security operation rather than keeping them siloed in a separate console?

The Capabilities Buyers Should Evaluate

Whatever framework a buyer uses to evaluate the market, the underlying capability set is now fairly consistent across serious vendors. The table below summarizes the capability areas security and risk management leaders should expect a modern email security platform to address, independent of any single vendor's marketing language.

Capability Area	What It Should Deliver
Core email protection	Baseline filtering for spam, known malware, and policy violations — table stakes, not a differentiator on its own.
Phishing & BEC detection	Behavioral, LLM- and NLP-based content analysis, including social-graph modeling of sender-recipient relationships, to catch impersonation, payment-redirect, and credential-harvesting attempts that lack a malicious payload — with the detection reasoning exposed to administrators, not just a pass/fail verdict.
Malicious links & attachments	URL and attachment analysis, including sandboxing, to catch payload-based threats before or immediately after delivery.
Content & artifact analysis	Deeper inspection of message content, headers, and embedded objects — including newer vectors like QR-code (“quishing”) payloads.
Threat intelligence	Use of external and internal indicators to distinguish emerging campaigns from one-off spam and to prioritize investigation.
User & behavioral risk	Identification of which individuals are repeatedly targeted, clicking, or exhibiting risky behavior — turning email data into a risk signal, not just a delivery log.
Investigation & incident response	A workflow for security teams to examine a suspicious message, its context, and related activity without switching tools.
Quarantine & remediation	The ability to act — pull a message from mailboxes org-wide, quarantine or disarm risky attachments, block a sender or domain, and document the action.
Outbound email security	Data loss prevention and encryption for sensitive outbound content, including financial data, credentials, and regulated data types.
SecOps / SIEM / XDR integration	The ability to feed email findings into broader detection and response workflows rather than operating as an isolated silo.



Two of these deserve a specific note, because they are easy to gloss over in a vendor pitch but are treated as baseline, named requirements in how the category itself is defined. First, behavioral and BEC detection should not be a black box: the platform should be able to show an administrator why a message was flagged which signals fired, not just a risk score since a verdict without reasoning is difficult to audit or defend to a board. Second, outbound protection is a data-protection function, not only a detection one: encryption of sensitive outbound mail belongs alongside DLP, not as a separate purchase.

Two further categories deserve emphasis because they are where the market is evolving fastest.

Enterprise vs. MSP/MSSP requirements diverge

A large enterprise typically wants deep integration with a single Microsoft 365 or Google Workspace tenant, tight alignment with its existing SIEM/XDR investment, and granular reporting for its own board and auditors. A managed service provider or MSSP, by contrast, needs the same detection depth applied across many customer tenants from one workflow, with the ability to triage and prioritize across those tenants rather than working one mailbox at a time. Buyers should ask vendors directly how multi-tenant operation is handled not assume that an enterprise-oriented tool scales cleanly to a services model, or vice versa.

AI is now an evaluation criterion, not a marketing bullet

Because generative AI has measurably improved attackers' phishing success rates, buyers should ask specifically how a vendor's AI is used: Is it applied to detection (identifying content and behavioral patterns of AI-generated lures), to investigation (summarizing and prioritizing alerts for analysts), or to both? Vague references to "AI-driven" protection without a specific description of what the model analyzes and how its output is used should be treated as a gap to probe in due diligence, not a checkbox to accept.

AI Is Changing Both Sides of the Fight

It is worth separating two distinct roles AI now plays in email security, because they are frequently conflated in vendor marketing.

AI as an attacker tool

Generative AI has removed the traditional tells of phishing poor grammar, awkward phrasing, generic greetings and made it economical for attackers to personalize lures at scale using publicly available information about a target's role, colleagues, and communication style. Microsoft's research found AI-automated phishing achieved roughly 4.5 times the click-through rate of traditional phishing attempts, and separately estimated that AI can make some phishing operations dramatically more profitable by allowing attackers to scale targeted campaigns with far less manual effort per target (Microsoft Digital Defense Report, 2025).

AI as a defensive capability

On the defensive side, machine learning models are now commonly used to classify message content and structure (distinguishing a genuine invoice request from a fraudulent one), to model sender-recipient relationships and flag anomalies (a first-time request for a wire transfer from an unfamiliar address claiming to be an executive), and to correlate weak signals across many messages and mailboxes that would be invisible when reviewing any single email in isolation.

For buyers, the useful question is not "do you use AI" nearly every vendor will say yes but "what specific decision does your model make, on what data, and how does a security analyst see and act on its output?" A platform that cannot answer that concretely is likely describing AI as a feature label rather than a functioning capability.

From Detection to Response: Why Workflow Matters as Much as Detection

A detection engine that flags a phishing email is only half the job. The other half often the half that determines whether an incident is contained in minutes or discovered weeks later during an audit is workflow: how a flagged message becomes an investigated incident, how a confirmed threat becomes an org-wide quarantine action, and how that action becomes a data point that improves the next detection.

Security teams evaluating email security platforms should look for four connected capabilities operating as one workflow rather than four separate features:

- **Visibility**: a consolidated view of message volume, inbound and outbound traffic, and overall email security posture, not just a list of blocked items.
- **Detection**: identification of phishing, BEC, malicious links, and suspicious activity using both content analysis and behavioral modeling.
- **Investigation**: the ability for an analyst to examine a flagged message, the user it targeted, related messages, and any connected security context, from a single workflow.
- **Response**: quarantine, notification, and the ability to route findings into broader security operations (a SIEM, a case-management tool, or a SOAR playbook) so that email-borne threats are handled with the same rigor as any other security event.

This is precisely where many legacy point tools fall short: they detect well but investigate and respond poorly, leaving security teams to manually reconstruct context who else received a similar message, whether the sender domain has been seen before, whether the same user has been targeted repeatedly using separate tools or manual log review.

As email increasingly functions as one stage in a longer, multi-vector attack chain (a phishing email that leads to a compromised credential that leads to lateral movement on the network), the ability to correlate email activity with identity and endpoint telemetry becomes a meaningful differentiator rather than a nice-to-have.

Seceon aiSecurity Email360: Applying These Principles

The preceding sections describe what the email security market broadly expects of a modern platform: behavioral detection, investigative context, and a direct path from detection to response connected to the rest of security operations rather than isolated in a standalone console. Seceon aiSecurity Email360 is built around that same logic, applied specifically to email telemetry.

A five-step workflow: Connect → Collect → Analyze → Correlate → Act

Rather than treating email security as a single filtering decision, Email360 structures it as a continuous workflow:

- **Connect:** Email360 connects to an organization's email environment and receives the activity needed to build a security view of the mail infrastructure.
- **Collect:** email activity is organized into structured security telemetry, covering messages, users, threats, and mailbox activity.
- **Analyze:** AI-driven analytics examine email behavior and threat indicators to identify suspicious activity and potential attacks.
- **Correlate:** events are correlated to reveal threat patterns, targeted users, incidents, and the broader security context behind email activity.
- **Act:** security teams investigate threats, manage quarantine, and route relevant findings into connected security workflows.

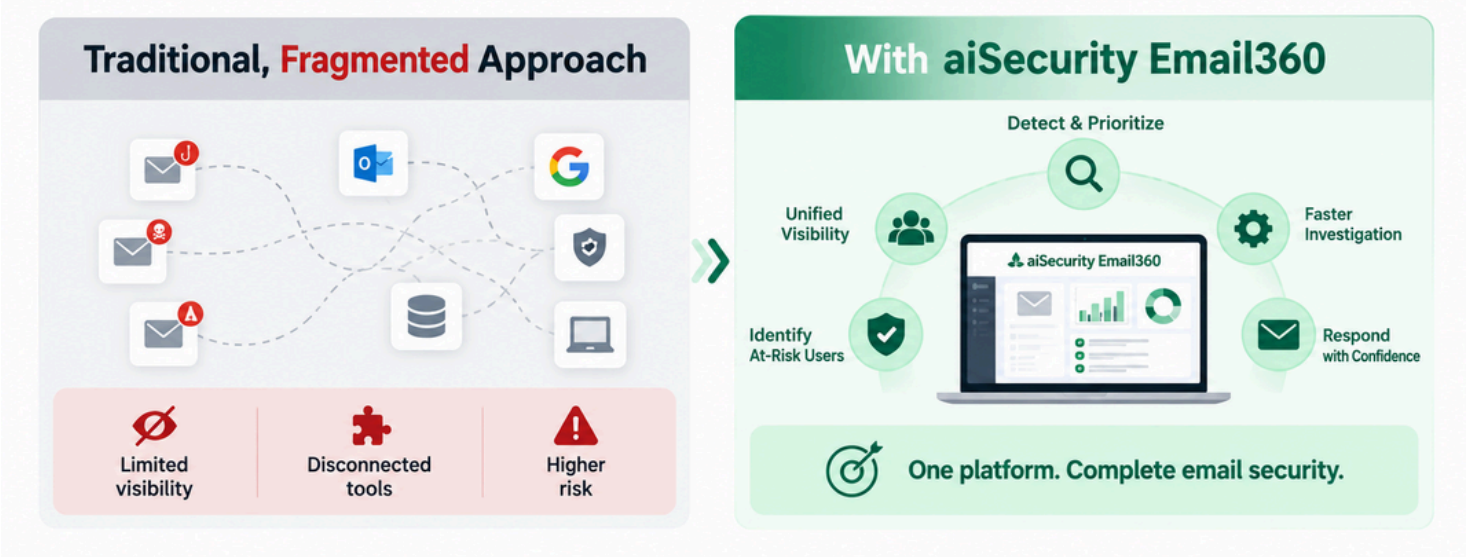
This mirrors the visibility-to-response lifecycle described in Section 5: Email360 is explicitly designed to move a security team from "here is what was blocked" toward "here is what is happening across our email environment, who is being targeted, and what we should do next."

What it detects and surfaces

According to Seceon's published product materials, Email360 is designed to detect phishing, business email compromise, malicious links, and other suspicious message activity using AI-driven security analytics, and to identify high-risk users and emerging attack patterns rather than evaluating each message in isolation. Suspicious messages, incidents, threat indicators, and targeted users are brought into a centralized investigation workflow, giving SOC teams the context needed to act quickly rather than piecing together evidence across separate tools.

Before and after: what changes for a security team

Traditional, Fragmented Approach	With aiSecurity Email360
Limited visibility into organization-wide email activity.	Unified visibility into email activity, threats, users, and security posture.
Threats and incidents scattered across multiple, disconnected tools.	A centralized view of suspicious messages, incidents, and email security events.
Difficult to identify which users are being targeted or are high-risk.	Identifies targeted users and highlights where email risk is concentrated.
Phishing, BEC, and malicious activity are hard to prioritize.	Helps detect and prioritize phishing, BEC, malicious links, and suspicious activity.
Security teams lack a unified view of email security posture.	Provides contextual intelligence for faster investigation.
Investigation and response workflows remain disconnected.	Connects investigation, quarantine, notifications, and broader security workflows.



Who Benefits, and How

Security teams and SOC analysts

For analysts, the practical value is a single workflow to detect phishing, BEC, and suspicious email activity, investigate incidents and targeted users, and quarantine or respond through connected security workflows rather than switching between an email gateway console, a ticketing tool, and a separate SIEM to reconstruct what happened.

CISOs and executives

For security leadership, Email360 is positioned to provide organization-wide visibility into email security posture and user risk, giving CISOs a defensible, evidence-based way to describe threat exposure to the board and to auditors a persistent gap independent breach and fraud data continues to point to, given how much financial and reputational risk continues to originate in the inbox.

MSPs and MSSPs

For managed service providers, the ability to deliver email security as a managed offering monitoring multiple customer environments from a unified workflow and identifying threats and high-risk users faster across that customer base addresses the multi-tenant requirement described in Section 3, where many point solutions built for a single enterprise tenant struggle to scale operationally.

Integration with the broader security operation

Because Email360 is part of Seceon's broader security operations platform, email findings are positioned to connect into the same investigation and response workflows used for other telemetry sources consistent with the SIEM/XDR integration capability buyers should look for in any modern email security investment, rather than treating email as a permanently separate silo.

Conclusion: Evaluating Your Next Email Security Investment

The data is consistent across independent sources: email remains the most heavily exploited entry point into the enterprise, the financial impact of business email compromise continues to climb, and generative AI is making social-engineering attacks measurably more effective. The category's shift toward layered, behavior-based, multi-vendor defense reflects a market that has outgrown the idea of email security as a single filtering decision made at the gateway.

Security and risk management leaders evaluating their next email security investment should look past headline detection-rate claims and ask harder questions: How does this platform detect behaviorally, not just by signature? How does it turn a flagged message into an investigated, documented incident? How does it identify which users are actually at risk, rather than only which messages were blocked? And how does it connect to rather than duplicate the rest of the security operation?

Those questions define the standard the market is converging on. Seceon aiSecurity Email360 is one vendor's answer to that standard, built around a connect-collect-analyze-correlate-act workflow intended to give security teams the visibility, detection, investigation, and response capabilities that email risk now demands as one integrated discipline within the SOC, not a separate mailbox chore.



Modern Email Security for a Visibility-Driven SOC

Detect. Investigate. Respond. Turn email into your strongest security signal.

The Email Threat Landcape

$3.05B
U.S. losses from Business Email Compromise in 2025

60%
confirmed breaches involve the human element

54%
click-through rate for AI-automated phishing

21%
of successful attack outcomes were BEC

Current Problems

- Limited real-time visibility into email activity
- Disconnected point solutions
- Growing volume of sophisticated attacks
- Manual investigation and response
- Difficult to identify targeted users
- Compliance and audit requirements
- AI-Generated Email Attacks
- Account Takeover & Credential Theft

Seceon Solution

- Unified, AI-driven email security platform
- Detects phishing, BEC, malicious links and attachments
- Identifies targeted users and risky behavior
- Automated response and quarantine
- Minimizes false positives
- Integrated with SIEM/XDR and broader SOC workflows
- Supports compliance and audit readiness

From Email Telemetry to Action

CONNECT → Integrate with your email environment

COLLECT → Ingest and normalize email telemetry

ANALYZE → Apply AI/ML to detect suspicious behavior

CORRELATE → Connect email with broader security context

ACT → Investigate, respond and quarantine

Quantifiable Security Outcome

95% reduction in false positives

Faster threat Detection

Improved SOC efficiency

Stronger compliance and audit readiness

Transform Email Security from Passive Filtering to Proactive Detection, Intelligent Investigation, and Faster Response.

About Seceon

Seceon enables MSPs, MSSPs, and Enterprises to reduce cyber threat risks and their security stack complexity while greatly improving their ability to detect and block threats, and breaches at scale. Seceon's Open Threat Management (OTM) platform augments and automates MSP and MSSP security services with our AI and ML-powered aiSIEM, and aiXDR platforms. The platform delivers gapless coverage by collecting telemetry from logs, identity management, network logs and flows, endpoints, clouds, and applications. It's all enriched and analyzed in real-time by applying threat intelligence, AI and ML models built on behavioral analysis, and correlation engines to create reliable, transparent detections and alerts. Over 850 partners are reselling and/or running the industry's lowest TCO, efficient security services with automated cyber threat remediation and continuous compliance for over 9,800 clients.



References and Citations:

This whitepaper is based on research and data from:

- Federal Bureau of Investigation, Internet Crime Complaint Center (IC3). 2025 Internet Crime Report. ic3.gov.
- Verizon. 2025 Data Breach Investigations Report (DBIR).
- Microsoft. Microsoft Digital Defense Report 2025. microsoft.com/security.
- Microsoft Threat Intelligence. "Email threat landscape: Q1 2026 trends and insights." Microsoft Security Blog.
- Seceon Inc. "aiSecurity Email360" product page. seceon.com/aisecurity-email360.

About the Author

Aditya Pandey

Software Engineer, Seceon Inc.



Aditya Kumar Pandey is a Software Engineer specializing in AI-driven email security, security orchestration, and automated remediation. His work addresses the evolving threat landscape by combining AI/ML, NLP, threat intelligence, and security analytics to identify and investigate sophisticated email attacks such as phishing and BEC. He designs automated security workflows that enable rapid containment, response, and remediation across enterprise security environments. His expertise spans SIEM, SOAR, Python, REST APIs, network security, cloud security, and security platform integrations. At Seceon, he focuses on transforming security intelligence into scalable, automated, and actionable cyber defense capabilities.

© 2026 Seceon Inc. All rights reserved. aiSIEM, aiXDR, are trademarks and registered trademarks of Seceon Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners.