The Essential Guide to Cloud Email Security

Best Practices for Securing Your Email Environment

Abnormal

Rethinking Email Security in the Age of AI

$2.77B

lost to business email compromise in 2024 alone.

FBI Internet Crime Report 2024

261 Days

is the average time required to identify and contain a phishing-related breach.

IBM Cost of a Data Breach Report 2024

84%

of employees fall for phishing emails within 10 minutes of receipt.

CISA Phishing Guidance 2023

Email remains the number one entry point for cyberattacks, but threats have advanced far beyond what most defenses were designed to detect. Today's AI-powered scams exploit identity, abuse trusted relationships, and weaponize everyday business communications. They originate from trusted accounts, reference real workflows, and unfold with surgical precision across your organization.

Attackers no longer rely on malware or suspicious links. Instead, they exploit trust: slipping into vendor conversations, hijacking executive mailboxes, and gaining access through authorized apps. These attacks blend into daily operations, and thanks to AI, they're happening faster and at a greater scale than ever before.

In this guide, we outline eight of the most pressing challenges facing cloud email security teams today, and the capabilities required to solve them. We also provide a ten-point checklist to support deeper evaluation and help ensure your next platform is truly built for the way attackers operate now.

When Attackers Have AI, Defenders Need It

The modern threat actor doesn't break in—they log in. With access to stolen credentials, MFA-bypass tools, and AI-generated social engineering kits, attackers exploit the most vulnerable layer in the organization: humans.

Phishing emails now replicate known brands and executive tone with alarming precision. Business email compromise (BEC) attacks follow up on real invoice threads, exploiting vendor relationships. And malicious third-party apps quietly siphon off email data after being approved by unsuspecting users.

These attacks succeed because they don't trigger legacy defenses like secure email gateways (SEGs) built to detect static threat indicators. That's why organizations need solutions that start with a model of what's normal in their daily operations and flag anything that deviates, using AI not as a bolt-on, but as the foundation of detection.

Cloud-Native Email Demands Cloud-Native Security

As threat vectors expand rapidly, fueled by the AI revolution, security teams can no longer rely on tools designed to filter yesterday's static attacks. What's needed now is a platform that understands how your organization communicates, flagging deviations in real time.

That's why AI-native platforms have become the new standard:

- They analyze behavior, not just content.
- They learn relationships, not just threat feeds.
- They remediate threats automatically, not hours later via triage queues.

Securing cloud email today means meeting AI-powered threats with AI-native defense.

Solving the 8 Most Pressing Email Security Challenges

Despite advances in cloud email infrastructure, many organizations still struggle to stop the full range of modern email threats. Traditional secure email gateways miss behavioral attacks that look legitimate, while some newer API-based solutions lack the depth to detect the newest compromise tactics. Although many organizations choose to replace legacy tools altogether, others use modern behavioral solutions to enhance the traditional protection they already have.

Regardless of the direction an organization chooses, email security today requires more than piecemeal features—it requires the right foundation. A truly modern platform must be cloud-native and API-first, designed to integrate deeply with Microsoft 365 and Google Workspace, adapt to changing workflows, and scale with evolving threats.

Still, what matters most is how that foundation translates into protection. The following challenges highlight the most critical gaps facing security teams today—and the capabilities required to close them. Each is paired with a modern solution rooted in behavioral AI, identity intelligence, and automation.

Challenge #1: Context-Aware Social Engineering Attacks Evade Traditional Detection

Modern attackers are moving beyond payload-based threats. Today's business email compromise, VIP impersonation, and phishing campaigns often rely on well-crafted, context-aware messages. With no links or attachments, they bypass traditional secure email gateways and rule-based filters entirely.

Generative AI is accelerating this shift, enabling attackers to automatically produce emails that replicate internal tone, mirror known workflows, and adapt to specific roles or relationships.

These attacks exploit trust, urgency, and routine business language, making them indistinguishable from legitimate communications. Static detection fails because these messages don't contain the payloads or known indicators that secure email gateways are built to find—leaving subtle, socially-engineered attacks to slip through unnoticed.

Key Issue:

Payload-free messages mimic trusted senders and real workflows, evading legacy defenses that focus on detecting known indicators of compromise.

Solution: AI-Native Protection Using Behavioral Intelligence to Stop Modern Social Engineering Tactics

The ideal detection engine analyzes every message against a behavioral baseline constructed uniquely for each organization. It ingests a wide range of signals—such as language patterns, tone, identity attributes, and communication history—to identify even the subtlest deviations from the norm.

AI enables this by correlating patterns across massive volumes of communication, learning what's normal for each user, department, and vendor, and spotting unusual behavior in ways a human or rule-based system would miss.

This context-rich approach enables precise detection of socially-engineered attacks that lack payloads or known indicators, blocking threats that evade traditional security without disrupting legitimate operations.

The Ideal Solution Will:

Stop context-aware attacks by modeling normal behavior and identifying subtle anomalies across language, tone, and communication patterns—without relying on outdated known threat indicators.

Challenge #2: Supply Chain Risk Goes Unmonitored

Vendor compromise is one of the most difficult threats to detect, especially when the sender appears legitimate. Attackers who hijack vendor accounts can insert themselves into invoice threads and request fraudulent payments, all under the cover of an established business relationship. These attackers now use AI to generate messages that mirror previous interactions, mimicking phrasing, cadence, and tone.

Despite this growing threat vector, many email security tools still treat all vendor messages the same, without analyzing communication behavior or flagging abnormal patterns.

Key Issue:

Without behavioral profiling, compromised vendors blend in with everyday business communication.

Solution: Behavioral Analysis for Vendor Communication

The ideal solution builds behavioral profiles for every vendor by monitoring communication cadence, message content, recipient patterns, and relationship history. AI enables this by analyzing thousands of signals across past conversations to determine what's normal for each vendor. When a vendor's behavior deviates—such as by sending requests at unusual times or introducing new recipients—the system should automatically flag or block the message for further review.

Unlike traditional tools that rely on static rules, AI makes it possible to detect these trust-based threats by learning each vendor's unique behavioral fingerprint.

This level of behavioral analysis allows security teams to effectively profile vendors, track communication history, and adapt to organizational changes over time—all through an automated system and with depth required to detect the most nuanced compromise attempts.

The Ideal Solution Will:

Protect against vendor compromise by analyzing behavioral patterns across supply chain communications, ideally through a federated database across all customers.

Challenge #3: Slow or Manual Threat Remediation Increases Risk

Security teams often rely on end-user reports or SOC queues to identify and remove malicious messages. This reactive model creates delays at every step—alerts must be submitted, investigated, and manually remediated. During that time, the email remains accessible, increasing the chance that a user will reply, forward, or click.

The challenge grows more urgent when threats blend in with legitimate traffic. Generative AI enables bad actors to easily create believable content that evades both filters and human intuition. Messages may mimic internal tone, reference familiar workflows, or appear to come from trusted sources. Even when a threat is eventually confirmed, the manual process means it may linger in multiple inboxes for hours, widening the window of exposure and increasing operational risk.

Key Issue:

Delayed manual response gives threats more time to spread or be acted upon by end users.

Solution: Automated, Instantaneous Remediation

The ideal solution continuously evaluates message risk and removes confirmed threats from all affected inboxes in real time, without waiting for manual triage or user reports.

AI plays a critical role here by scoring risk based on behavioral signals, language cues, and sender reputation—automatically identifying threats as they emerge and initiating remediation without human input.

Automation is especially critical when threats mimic legitimate communications or originate from trusted accounts. Fast, reliable remediation prevents these messages from lingering in inboxes and limits the window of exposure across the organization. It also reduces the time SOC teams spend on investigation and reporting, freeing up analysts to focus on high-impact security initiatives.

The Ideal Solution Will:

Remediate threats instantly across all inboxes, minimizing user exposure and drastically reducing the burden on security teams.

▸▸ Challenge #4: Internal Email Threats from Hijacked Accounts Go Undetected

When attackers gain access to employee inboxes, they often use those accounts to launch phishing attacks internally, posing as trusted colleagues. These lateral attacks are hard to catch when email security solutions focus only on monitoring inbound east-west messages.

Increasingly, these attackers leverage AI to automate internal compromise and generate tailored messages that reflect the language, tone, and timing of real employee communications.

Without visibility into internal behavior, compromised accounts can be used to request sensitive data, redirect payments, or escalate access—all without triggering alerts.

Key Issue:

Internal account misuse often appears normal, bypassing tools that focus solely on external threats.

Solution: Lateral Threat Detection with Identity and Behavioral Context

The ideal solution builds behavioral baselines for every employee by analyzing login activity, device usage, and internal communication patterns. When an account begins behaving abnormally—sending unusual messages, accessing sensitive data, or logging in from unfamiliar locations—the system should automatically flag and respond to the threat.

AI enables this by learning each user's normal behavior across communication and identity signals. Effective detection requires visibility into east-west email traffic, not just inbound threats. By continuously monitoring for deviations in user patterns and internal message flow, organizations can detect and contain compromised accounts before they're used to escalate attacks or exfiltrate data.

The Ideal Solution Will:

Detect and contain internal threats by monitoring identity signals and behavioral deviations across internal email activity—stopping lateral movement before it causes harm.

Challenge #5: SOC Teams Miss Early Signs of Account Takeover

Bad actors don't need to steal passwords when they can bypass authentication altogether. By exploiting legacy login protocols, phishing for session tokens, or embedding malicious apps, they quietly gain persistent access to employee accounts, often without raising alerts.

With the help of AI, attackers can easily identify these vulnerable entry points, mimic legitimate login behavior, and escalate access across environments—making their activity look routine and going undetected until it's too late.

Legacy tools may flag suspicious activity once the attacker acts, but rarely identify the early signs of compromise. By then, it's too late.

Key Issue:

Without visibility into identity signals and app behavior, SOC teams can't stop takeovers before damage is done.

Solution: Email Account Takeover Protection with Automatic Remediation

The ideal solution monitors sign-in activity across all users, analyzing factors like geographic location, device type, login patterns, and authentication methods. It should identify anomalies such as impossible travel, unfamiliar browsers, or suspicious access attempts that may indicate a compromised account.

AI is key to correlating these identity signals in real time, learning what's normal for each user and surfacing risk when login behavior, device access, or app usage deviates in meaningful ways. This isn't bolted-on AI for marketing purposes. It's core detection logic that processes thousands of behavioral signals continuously to make precise, real-time decisions without human intervention.

In addition to login activity, the solution must also surface risky third-party applications—especially those with excessive permissions or abnormal behavior—before they can be used to escalate access or exfiltrate data.

The Ideal Solution Will:

Detect and respond to account takeovers early by monitoring identity signals and application behavior—preventing escalation and lateral movement.

Challenge #6: Lack of Visibility Into Email Platform Configurations and App Integrations

Many email security tools stop at the inbox. They don't monitor tenant-level configuration changes, permissions granted to apps, or identity drift across the cloud environment. That leaves security teams blind to risks introduced through misconfiguration or abuse of trust.

Attackers routinely exploit these blind spots—gaining access not by hacking, but by inheriting permissions no one knew were risky. With the help of AI, they can map misconfigured environments, target over-permissioned apps, and blend malicious activity into routine admin changes, making threats harder to spot and quicker to escalate.

Key Issue:

Misconfigurations and over-permissioned apps go unnoticed, opening hidden pathways for attackers.

Solution: Security Posture Management with Continuous Monitoring

The ideal solution provides continuous visibility into changes across users, applications, and configurations within cloud email platforms like Microsoft 365 and Google Workspace. It should detect signs of configuration drift, unsafe privilege escalations, and unauthorized third-party app integrations.

AI enhances this capability by monitoring and interpreting large volumes of activity and configuration data, flagging patterns that deviate from established norms and surfacing potential exposure before it becomes exploitation.

By proactively monitoring the broader email environment—not just the inbox—security teams can identify and address misconfigurations or trust abuses before they turn into breaches. This visibility is essential for maintaining a strong security posture in dynamic, cloud-first environments.

The Ideal Solution Will:

Surface configuration risks and unsafe app permissions across the email ecosystem, enabling proactive hardening and reducing the risk of silent compromise.

Challenge #7: Graymail and Spam Degrade Productivity and Focus

Beyond attacks, executives and employees alike lose time to unnecessary email: newsletters, promotions, and bulk senders that distract from critical work. This noise doesn't just affect productivity—it can also camouflage real threats and facilitate social engineering attacks like email bombing.

Security tools that rely on manual rules or quarantine filters force users to manage the problem themselves, creating friction and frustration.

Key Issue:

Time-wasting email reduces focus, increases risk, and burdens end users with unnecessary sorting.

Solution: AI-Powered Graymail Management and Productivity Insights

The ideal solution uses behavioral AI to understand user preferences and automatically classify graymail without relying on static filters or manual rules. By deprioritizing or removing unwanted messages based on engagement patterns, such a system helps keep inboxes focused and relevant.

In addition to improving the user experience, the solution should provide insights into the time saved and overall productivity impact, giving IT and security leaders a clear view of the business value of a cleaner inbox.

The Ideal Solution Will:

Reduce graymail and spam through behavioral filtering, boosting productivity while maintaining insight into communication trends.

▸▸ Challenge #8: Traditional Email Security Awareness Training Falls Short

While preventing malicious emails from ever reaching the end user should always be the goal, no solution is 100% effective all the time. As a result, organizations must train their users to recognize the signs of attack and respond appropriately.

Unfortunately, most security awareness programs take a one-size-fits-all approach, delivering static, outdated training modules that feel irrelevant to many employees. These programs often lack context, fail to reflect the real threats users face, and rely on infrequent phishing simulations that don't match modern attack techniques.

Worse, traditional training is time-consuming to manage and expensive to scale. It places a heavy burden on security teams to create content, schedule campaigns, and track results, often with little evidence that the training is improving employee judgment or reducing risk.

Key Issue:

Traditional awareness training is static, generic, and disconnected from the real threats employees encounter.

Solution: Personalized, Contextual Security Awareness Training

The ideal solution delivers personalized training experiences that reflect each employee's role, exposure to specific threats, and behavioral patterns.

AI is essential to this approach. By analyzing real attack data and user interactions, AI-powered training adapts to provide relevant simulations and timely feedback. This ensures that employees are engaged with content that mirrors actual threats they may encounter, enhancing their ability to recognize and respond to phishing attempts. It also reduces administrative overhead by automating the delivery and adjustment of training modules based on evolving risks and individual behavior.

The Ideal Solution Will:

Provide personalized, adaptive training that reflects real-world threats and user behavior—improving engagement, strengthening judgment, and reducing the burden on security teams.

An Abnormal Solution to Cloud Email Security

Abnormal AI is the only cloud email security platform built from the ground up to detect and stop socially-engineered attacks with behavioral AI. While traditional tools rely on static rules and known indicators, Abnormal models the normal, analyzing thousands of signals across identity, content, and communication patterns to detect what others miss.

With one-click API integration, Abnormal connects directly to Microsoft 365 or Google Workspace and begins building a behavioral baseline unique to each organization. The platform evaluates every message, login, and configuration change in real time—enabling precise detection, instant remediation, and comprehensive visibility across the email environment.

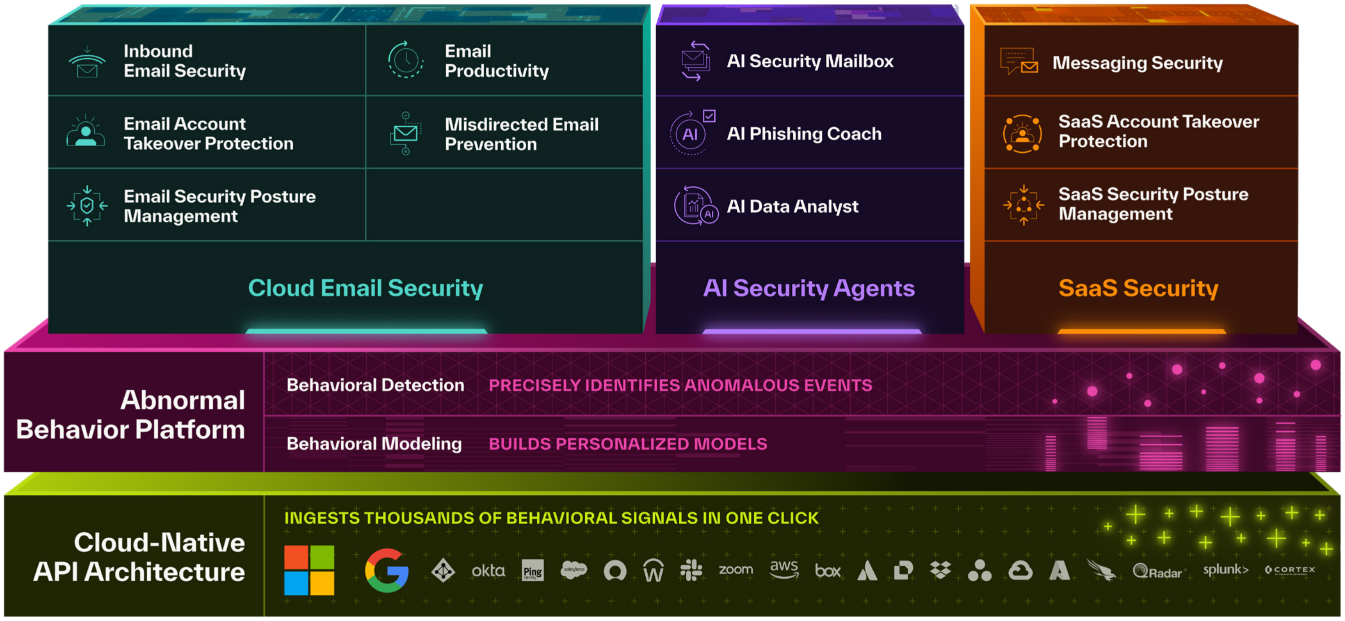
Inbound Email Security
Email Productivity
Email Account Takeover Protection
Misdirected Email Prevention
Email Security Posture Management
Cloud Email Security
AI Security Mailbox
AI Phishing Coach
AI Data Analyst
AI Security Agents
Messaging Security
SaaS Account Takeover Protection
SaaS Security Posture Management
SaaS Security
Abnormal Behavior Platform
Behavioral Detection
PRECISELY IDENTIFIES ANOMALOUS EVENTS
Behavioral Modeling
BUILDS PERSONALIZED MODELS
Cloud-Native API Architecture
INGESTS THOUSANDS OF BEHAVIORAL SIGNALS IN ONE CLICK
okta
Ping
zoom
aws
box
Radar
splunk>
CORTEX

Abnormal Platform Solutions

Each solution in the Abnormal platform addresses a specific area of risk highlighted in this guide, helping organizations detect advanced threats, automate incident response, and improve security outcomes across email and collaboration.

Inbound Email Security

Leverages behavioral AI to detect and block advanced email threats—such as business email compromise, credential phishing, and supply chain attacks—that traditional solutions like SEGs miss. By understanding normal communication patterns unique to each organization, Abnormal identifies anomalies and prevents malicious emails from reaching inboxes.

Email Account Takeover Protection

Monitors user behavior and sign-in activity to detect compromised accounts in real time. Automatically responds by disabling access, terminating sessions, and initiating password resets to prevent further unauthorized activity.

AI Security Mailbox

Automates the analysis and remediation of user-reported phishing emails. Reduces SOC workload by quickly identifying threats and providing employees with timely feedback on their reports.

Email Security Posture Management

Continuously assesses the email environment for misconfigurations and risky settings. Provides actionable insights to remediate issues that could be exploited by attackers, enhancing overall security posture.

Email Productivity

Utilizes AI to learn individual user preferences and filter out graymail—such as newsletters and promotional emails—reducing inbox clutter and improving focus. Provides analytics on engagement and time saved.

AI Phishing Coach

Turns real phishing attacks into personalized simulations for each employee. Delivers just-in-time training and feedback, reinforcing good security habits and reducing susceptibility to future attacks.

With a full suite of AI-native solutions, Abnormal brings together detection, response, posture management, and user resilience, powered by behavioral data to protect the full spectrum of cloud-based email and collaboration.

This unified approach eliminates silos, reduces operational overhead, and delivers the precision modern security teams need to stop advanced threats before damage is done. That's why Abnormal is trusted by more than 3,200 organizations, including over 20% of the Fortune 500.

Selecting the Right Email Security Platform

Modern cloud email attacks no longer rely on payloads or perimeter gaps—they exploit behavior, trust, and identity. This guide has outlined the new challenges that many platforms struggle to address, and the capabilities required to meet them. As you evaluate solutions, the right questions will reveal whether a platform is truly equipped to stop AI-powered attacks, support your team, and scale with your environment.

The next generation of email security doesn't just block threats. It understands how your organization works and adapts to protect it.

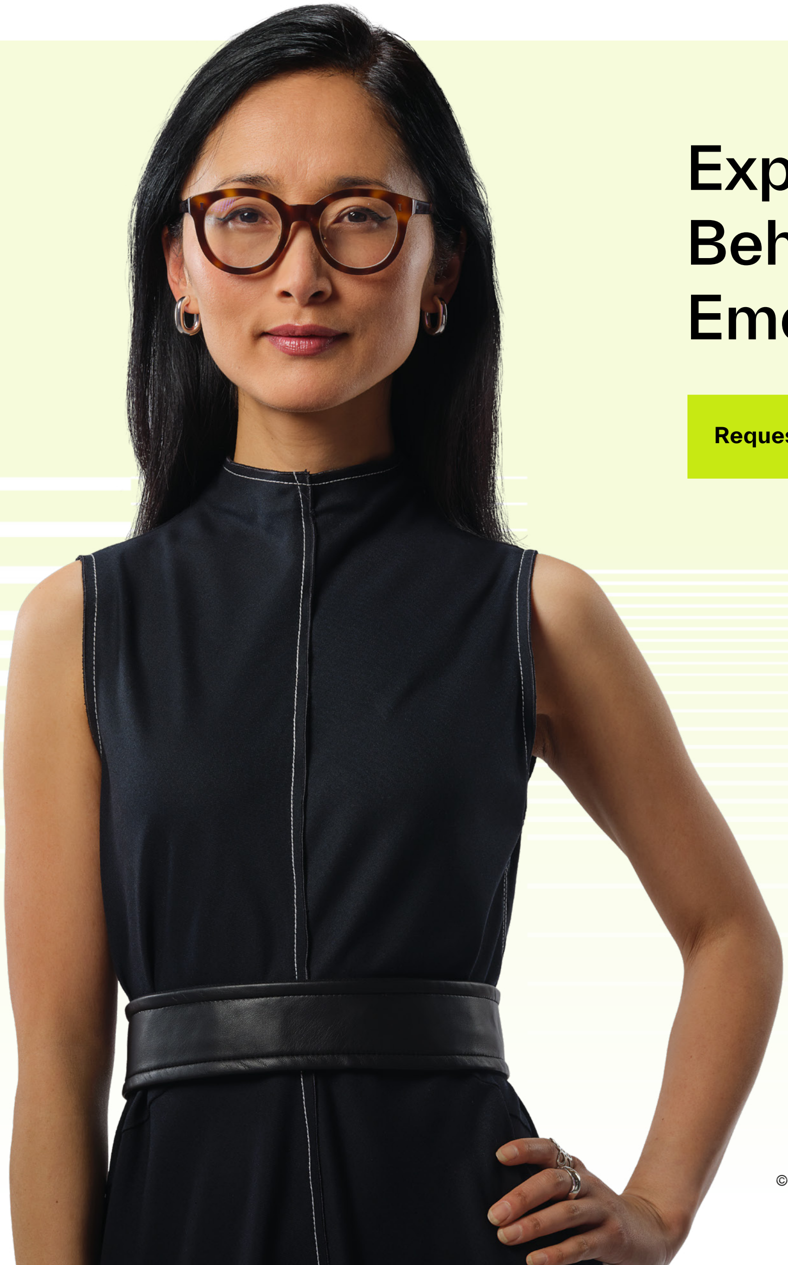

Experience an AI-Native Behavioral Approach to Email Security

Request a Demo >

Read Customer Stories >

10 Must-Ask Questions for Choosing the Best Cloud Email Security Solution

Choosing an email security platform isn't just about ticking boxes. It's about ensuring the technology can address the ever-shifting realities of today's threat environment. In addition to our list of most important challenges, we've designed these 10 questions to support a more informed evaluation: to uncover meaningful differences between vendors, move beyond surface-level claims, and focus on what truly matters—detection accuracy, behavioral context, automation, and user impact. Use this checklist to ensure you pick the email security platform best suited for your organization.

1. Can the solution detect attacks that don't include payloads or known indicators?

Legacy tools rely on scanning for malware, bad links, or known domains. Today's threats often include no payload at all—just persuasive language and human trust.

- □ Does the platform use behavioral AI to analyze tone, content, and sender context?
- □ Can it detect payload-free BEC, impersonation, and VIP fraud?
- □ Will it identify threats that look normal but act abnormally?

2. Does the platform natively understand your employees, vendors, and their behavior?

Effective detection starts with knowing what's normal. A modern platform should go beyond mere scanning of anomalies in content; it should build consistent baselines across internal and external communications.

- □ Can the solution profile users, vendors, and communication history?
- □ Does it understand invoice frequency, typical recipients, and message style?
- □ Does it adapt to organizational changes over time?

3. Can it detect internal threats like lateral phishing or insider compromise?

Once an attacker gains access to a legitimate account, they can launch phishing campaigns from inside your environment, bypassing external defenses and abusing trust.

- □ Does the solution monitor east-west (internal) traffic, not just inbound?
- □ Can it detect abnormal behavior from real employee accounts?
- □ Does it use consistent but flexible behavioral baselines to flag internal misuse?

4. Is threat remediation instantaneous, or delayed by manual workflows?

Even a short delay increases the chance that an employee clicks, replies, or forwards a malicious email. The longer threats linger, the higher the risk—and the heavier the burden on SOC teams. Manual triage and investigation consume valuable analyst hours, especially when dealing with user-reported phishing and abuse mailbox traffic.

- □ Does the platform remediate threats in milliseconds, not minutes or hours?
- □ Can it remove messages from every affected inbox automatically?
- □ Does it reduce SOC workload by automating triage and surfacing only high-risk messages?

5. Can it automatically detect and stop account takeovers—before internal attacks happen?

Attackers easily exploit legacy authorization, MFA fatigue, and malicious apps to compromise accounts without detection.

- □ Does it monitor sign-in activity, device type, and login patterns?
- □ Will it flag MFA bypass or risky app installs in real time?
- □ Can it take action before compromised accounts are used for lateral phishing?

6. Does the platform give visibility into configuration risks and cloud app permissions?

Misconfigurations, overly permissive apps, and tenant-level changes are some of the easiest ways for attackers to gain footholds. Yet, many email security tools don't properly monitor them.

- □ Can the platform surface risky changes in Microsoft 365 or Google Workspace?
- □ Does it detect over-permissioned or suspicious third-party apps?
- □ Will it alert on tenant posture drift before leading to exposure?

7. How much value does the solution deliver from AI—vs. using it for trendy marketing?

AI is everywhere now, but not every implementation is meaningful. Bolted-on AI rarely delivers the depth or precision needed to detect modern threats.

- □ Is AI central to how threats are detected, scored, and remediated?
- □ Does the platform learn from user and organization-specific behavior?
- □ Can it explain its decisions in a way analysts can trust?

8. Does it reduce end-user noise while keeping employees productive?

Too many alerts, irrelevant emails, and digest fatigue erode attention and hide real threats. Email security should clean up the inbox, not clutter it.

- □ Can the solution suppress graymail based on behavior and engagement?
- □ Does it reduce time spent managing quarantines and digests?
- □ Can it quantify productivity gains for different user groups?

9. Does the platform provide relevant, real-time security awareness training?

Security awareness training matters—but too often, programs rely on static, one-size-fits-all content that fails to reflect modern threats. Many simulations are outdated, disconnected from real attack patterns, and time-consuming to manage. The result: low impact, high effort, and limited improvement in user behavior.

- □ Is the training customized to reflect the specific risks each employee faces, rather than relying on outdated, one-size-fits-all content?
- □ Can the platform convert real attacks into timely simulations that deliver training when it's most impactful?
- □ Can security teams measure impact without spending hours managing content and campaigns?

10. Is it built for the future—or just retrofitted for today?

Some solutions claim to be cloud-native but were originally built for a different era. Ask how deeply the architecture supports modern workflows, with an eye toward future re-configurations and threat vectors.

- □ Was the platform built API-first, or adapted from SEG infrastructure?
- □ Can it deploy quickly and integrate deeply with cloud providers and pre-existing security tools?
- □ Does it protect communication platforms across the SaaS environment, beyond email?