ANALYST REPORT

2025 State of Misdirected Email Prevention

Keeping Sensitive Data Out of the Wrong Inboxes

Abnormal

Executive Summary

When CISOs hear "email threat," they imagine phishing attacks, credential harvesting, or business email compromise, not an employee accidentally sending confidential data to the wrong "Tom." Misdirected emails appear legitimate because they *are* legitimate, but these mistakes nonetheless result in data breaches, compliance violations, and costly remediation for organizations across industries. According to Gartner, human error is the number one cause of email-vectored data leakage.

Traditional email security solutions weren't designed to stop messages from reaching unintended recipients. Their primary focus on inbound threats and rules-based detection leaves these solutions unequipped to address misdirected emails that occur as part of normal outbound communication. Secure email gateways (SEGs) and traditional data loss prevention (DLP) tools lack the behavioral context needed to understand why sending *this particular email* to *this particular recipient* carries risk.

Misdirected email prevention is an aspect of enterprise cybersecurity that's long been neglected, and traditional tools made it difficult to address. Today's behavioral AI solutions change that. Because these solutions establish a baseline of normal communication patterns, they can flag incidents that stand out. Behavioral AI can read signals in email content, login locations, and sender-recipient relationships to uncover underlying intent.

This gives modern solutions access to a layer of insight that traditional tools lack, making it possible to:

- **Detect misdirected email accurately**
- **Stop accidental data loss automatically**
- **Accelerate remediation**
- **Simplify compliance**

Most enterprises have yet to incorporate behavioral AI into their email security stack. But as cybersecurity becomes increasingly AI-driven, it's clear that the status quo is no longer effective at reducing the risk of misdirected emails.

To better understand their progress in implementing modern AI-powered safeguards against misdirected email, we surveyed more than 300 security and IT stakeholders in organizations ranging in size from 1,000 employees to over 25,000. Participants hold a variety of positions, with 22% serving in leadership roles including CIO, CTO, CISO, or VP of IT or of Security. More than two-thirds (69%) are directors, managers, or team leads in security operations (SecOps).

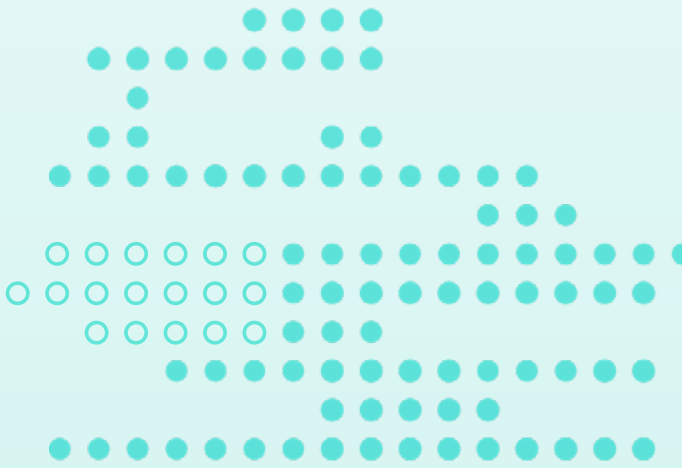

In this report, we explore our findings. Readers will learn what security and technology professionals are currently doing to prevent accidental data loss through misdirected email (and how they're thinking about human error-related cyber risk at large). We'll also explore the limitations of current solutions, and consider what's needed to solve the misdirected email problem proactively.

96%

of survey participants reported that their organization had experienced data loss or exposure due to misdirected email within the past year.

95%

had seen a measurable business impact (such as remediation expenses or a compliance violation) from these incidents.

98%

of respondents believe that misdirected email represents a significant risk to their organization.

41%

of organizations most often learn about misdirected emails when they're reported by the recipient.

69%

of survey participants are looking for a solution that can accurately and automatically block misdirected emails before they're sent.

97%

favor leveraging behavioral AI in a misdirected email prevention solution.

Table of Contents

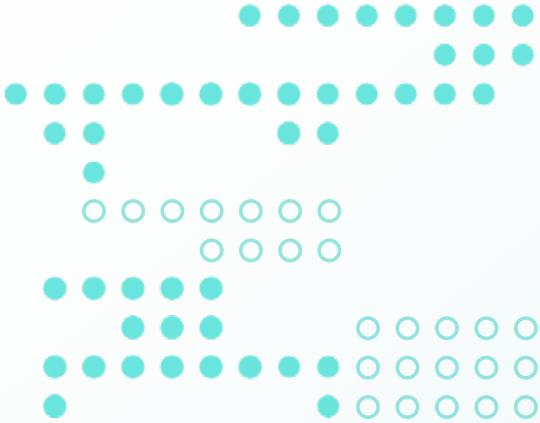

Misdirected Email: The Business Impact Is Real, and It's Costing Millions 5

- Misdirected Email Risks Are Often Overlooked 6
- Misdirected Email Is a Leading Cause of Data Loss 7
- A Single Mistake Is One Too Many 9

The Visibility Gap: Why the Problem Persists 10

- How Organizations Spot Misdirected Emails 11
- Capability Gaps in Today's Tech 12
- False Positives Drain Resources 13

A New Formula for Success 14

- Improving Misdirected Email Protection 15
- The Status Quo Is Not Working 16
- It's Time for Something New 17
- Behavioral AI Is the Answer 18
- Abnormal Misdirected Email Prevention 20

Conclusion 21

About Abnormal 22

Misdirected Email: The Business Impact Is Real, and It's Costing Millions

Security leaders often focus on external attacks and malicious data exfiltration, yet one of the leading causes of data loss is neither malicious nor intentional: misdirected email. According to the 2025 Verizon Data Breach Investigations Report, nearly 72% of breaches caused by end-user actions involve email misdirection, making this a top cause of data loss across industries. In fact, Gartner states that human error is the number one cause of email-related data leakage.

These might seem like small mistakes, but their consequences are outsized. Misdirected email can result in intellectual property loss, reputational damage, and compliance violations. In 2024, misdirected emails were the most frequently reported GDPR violation, representing a 27% year-over-year increase and a sizable chunk of the $1.2 billion in fines that were levied.

Misdirected Email Risks Are Often Overlooked

Our survey participants are experiencing significant numbers of accidents involving human error, sensitive data, and email. The vast majority (96%) reported that their organization had encountered data loss or exposure due to misdirected email within the past year. Over one-third of respondents (35%) had experienced data compromise due to misdirected email 10 or more times within that timeframe, and 13% had experienced it 20+ times.

Within organizations observing large numbers of misdirected email incidents, security executives and managers are more likely to notice the problem than IT leaders. Only 4% of IT executives reported that their organizations had detected more than 20 misdirected emails within the past year, whereas 16% of security leaders and 14% of security managers reported the same issue. This may indicate greater awareness—or simply better visibility—among participants in security roles.

Still, this risk often receives insufficient attention and often falls within a blind spot in security coverage. Legacy email security solutions (such as SEGs) tend to focus solely on inbound threats, and this is often where security teams have the most visibility. But inbound threats are only one piece of the puzzle. While DLP covers a broad range of outbound risk, significant blind spots in misdirected email coverage remain. To mitigate all email-related business risks, leaders must adopt a more holistic approach.

How many times has misdirected email resulted in data loss/exposure for your organization in the past year?

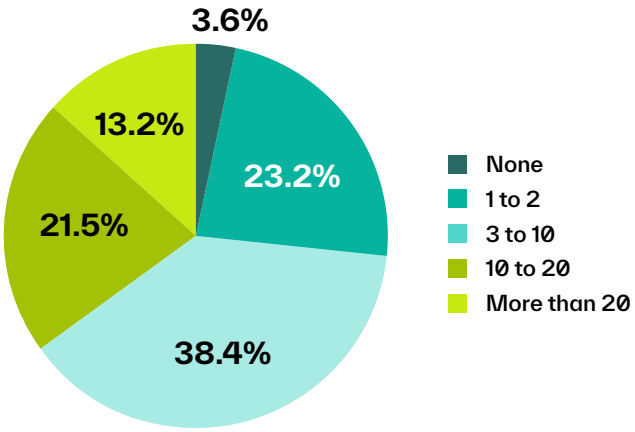

What Is Misdirected Email?

Misdirected email occurs when a message is sent to an unintended recipient. The misdirection could be caused by human error, a technical mishap, or another accident. Misdirected email is often thought of as a minor mistake, but these messages may contain sensitive information such as customer or financial data, intellectual property, or confidential business discussions. In these circumstances, misdirected email creates security and compliance risks—and constitutes a data breach.

In a world where most data loss takes place via email, misdirection represents a silent but serious breach risk—one that security leaders can no longer afford to ignore.

Misdirected Email Is a Leading Cause of Data Loss

Cybersecurity stakeholders know that email-related data loss represents a major business risk, and they're aware that misdirected email is responsible for a significant portion of these incidents. It's no surprise that survey participants rank misdirected email among the top data security risks their organizations currently face.

A full 98% of respondents rated misdirected email risks as "significant" in comparison to other threats to their data's confidentiality and integrity—such as malware or deliberate exfiltration. Two-thirds (67%) consider misdirected email an "extremely significant" or "very significant" risk to their organizations. Security executives are particularly attentive to these risks, with 90% of them rating misdirected email an "extremely" or "very" significant threat.

The more misdirected email incidents an organization experiences (or notices), the more likely its employees are to consider misdirected email a major risk. Among respondents in organizations seeing more than 20 instances of misdirected email each year, 45% consider this an "extremely significant" risk, while 33% see it as "very significant." Among those experiencing between 11 and 20 incidents per year, 34% view misdirected email as an "extremely significant" risk and 39% consider it "very significant."

These concerns reflect widespread awareness among stakeholders that misdirected email is ubiquitous. Its causes lie at the intersection of human nature and typical business processes—that is, the ways we are used to working. Because it's embedded in our habits, training or awareness programs alone cannot solve this problem.

In your view, how significant is misdirected email compared to other data loss risks your organization faces (e.g., malware, insider threats, data exfiltration)?

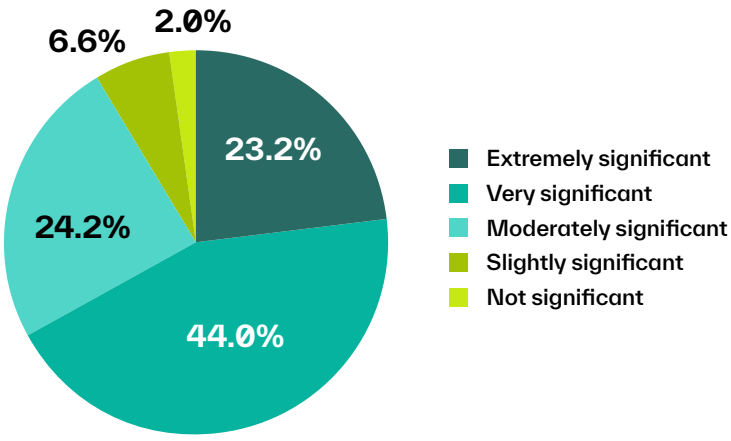

What Causes Misdirected Email?

The vast majority of misdirected emails result from simple, common mistakes.

- **Typographical errors:** A single misplaced letter or minor misspelling can send sensitive information to the wrong recipient.
- **Autocomplete mistakes:** Email programs often suggest the wrong contact based on recent communications, and quick clicks can result in error.
- **Similar-looking addresses:** It's easy to confuse recipients with similar names, especially in large organizations.
- **Workplace fatigue and time pressure:** Employees tasked with managing large volumes of email are more likely to make mistakes, especially when racing to meet deadlines or working under distraction.

If your organization relies on certain systems or workflows, it's likely that your misdirected email risks are higher than average. These include:

- **Outdated or overly broad distribution lists:** Old email groups may include individuals who should no longer be receiving sensitive communications.
- **Ineffective autofill logic:** Most email programs prioritize frequency of contact rather than more nuanced contextual information, leading them to propose recipients who may not be authorized to view the contents of the message.

Human error, organizational inefficiencies, and technical limitations are at fault in most cases of email misdelivery.

A Single Mistake Is One Too Many

Misdirected email is a major source of data loss, leading to remediation expenses, regulatory penalties, and erosion of trust. More than nine in ten survey participants (95%) work for an organization that experienced a serious business impact from misdirected email within the past year.

These impacts are both widespread and far-reaching. More than half (54%) of respondents' organizations had to spend large amounts of time, money, or effort remediating the incident. Nearly as many (49%) suffered the loss of confidential data in the incident. And an alarmingly large minority (40%) saw damage to their customer relationships because of misdirected email.

The number of misdirected emails that an organization observed did not always correlate with the extent of the damage they caused. Organizations that experienced 20 or more misdirected email incidents did report more impacts across all categories than those that had seen only one or two. Still, 51% of the organizations that had observed one to two misdirected email incidents saw confidential data exposed in those incidents, whereas only 46% of those with 11 to 20 incidents reported the leakage of confidential data. And 50% of those that had observed only one to two incidents reported experiencing excessive remediation costs or effort, while 53% of those that had experienced 20+ incidents reported the same thing. This is only a slight difference in business outcome, despite a substantial difference in the number of misdirected emails identified.

Regulatory Exposure

Multiple regulatory regimes across industries and geographies treat a single misdirected email incident as a compliance violation. These include:

- **Europe's General Data Protection Regulation (GDPR):** Any personal data sent to an unintended recipient qualifies as a data breach. Organizations subject to GDPR must report breaches within 72 hours and may incur fines of up to €20 million or 4% of annual global revenue.
- **The Health Insurance Portability and Accountability Act (HIPAA):** Misdirected emails containing protected health information (PHI) require breach notifications and can result in significant financial penalties.
- **Financial regulations:** FINRA rules and the provisions of the Sarbanes-Oxley Act (SOX) explicitly prohibit the disclosure of confidential financial data to unauthorized parties, regardless of whether that disclosure was intentional.

These discrepancies underscore an important point: It only takes one misdirected email incident to cause significant financial, reputational, or operational damage.

Which impacts has your organization experienced within the past year as a result of misdirected emails? Select all that apply.

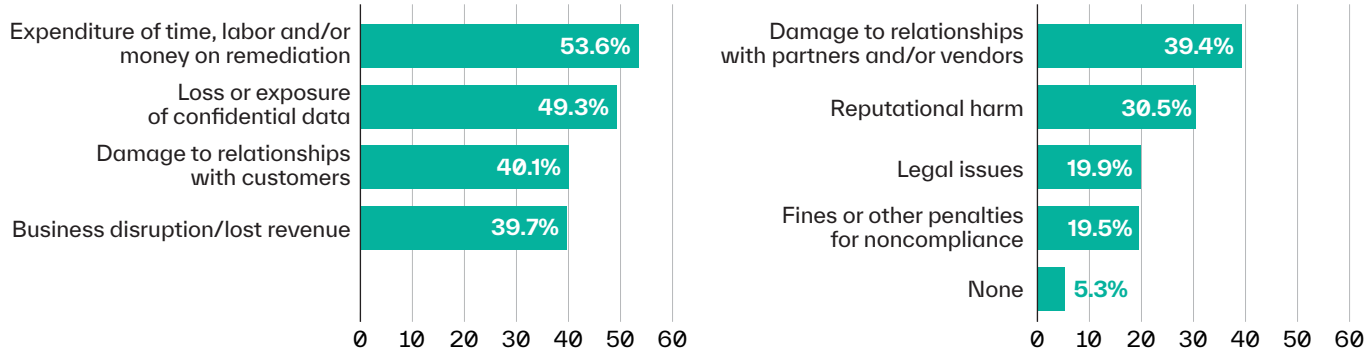

The Visibility Gap: Why the Problem Persists

Organizations can't stop what they can't see. All estimates about the extent of the misdirected email problem are just that—estimates—because incidents that go undetected are left out of the calculation. Recipients are under no obligation to report misdirected email, and may not notice that the communication was meant for someone else. Senders aren't always aware of their mistake. Traditional DLP and email security solutions (e.g., SEGs) rely on rules-based detections that weren't designed to catch things like one client's deliverables accidentally being sent to a different client.

Though these solutions aren't effective at detecting or proactively blocking misdirected email, they *are* noisy. Security teams often spend dozens of hours investigating the false positive alerts they generate without achieving the desired result: visibility and control.

How Organizations Spot Misdirected Emails

Participants in our survey know firsthand that proactive detection of misdirected email isn't easy. One of the most common ways misdirected email is identified is through voluntary reporting by the recipient. Though this was the third most common method by which teams learned about misdirected email incidents (mentioned by 41% of survey participants), it's worth noting that recipient reporting implies that existing email security solutions, DLP tools, or other security safeguards were unable to prevent the incident. By the time the recipient reports the incident, the damage (such as a HIPAA or GDPR violation, or breach of confidentiality) has already occurred.

Our findings also show that existing tools and solutions are missing the mark nearly half of the time. For instance, 47% of organizations don't typically learn about misdirected emails from their email security solution, and 49% aren't alerted by their DLP tool. Email security solutions are the number one detection method for misdirected email, but significant improvements are still needed, especially among some of the industry's most widely deployed technologies.

Which are the most common ways that your organization learns about misdirected email? Select three.

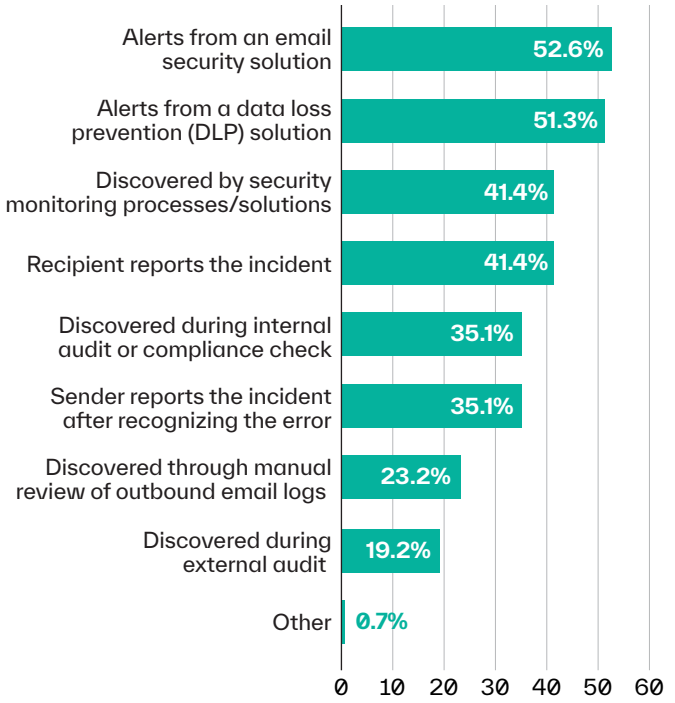

Why Traditional Security Tools Fall Short

Conventional data protection solutions weren't designed to catch employee mistakes, and legacy email security tools weren't built with a focus on outbound email. These solutions tend to be limited by:

- **Rigid dependence on rules:** Traditional DLP tools look for predefined keywords or patterns within data. They're not intended to identify human error, which requires a contextual understanding of the communication.
- **The tendency to assume internal communications are safe:** Misdirected email sent within the organization usually won't trigger alerts because traditional tools assume that all intra-domain communications are safe.
- **A lack of behavioral understanding:** These systems cannot recognize anomalous behavior, such as when an employee sends financial data to an unusual recipient.

Capability Gaps in Today's Tech

Even when a robust email security solution is in place, preventing misdirected email requires an approach grounded in an understanding of user behavior and communication context, not just static rules. Few vendors offer this.

When we asked survey participants about their top challenges in preventing data loss through misdirected email, we learned that they tend to focus on policy creation and enforcement. These capabilities are important, but they're not what's required to identify misdirected email before the message is sent.

The most frequently cited inhibitor to preventing data loss through misdirected email was difficulty with policy enforcement across complex hybrid environments (59% of respondents). However, even when accurately enforcing granular policies, most solutions cannot catch these mistakes. After all, these are legitimate communications, containing information that's permissible to share—except they have been sent to the wrong recipient. Without AI and the nuanced understanding of communication patterns and context that it enables, misdirected email is nearly impossible to detect.

Another top-ranked challenge was creating clear, consistent policies for safely sharing sensitive data via email. This was mentioned by slightly more than half (52%) of survey participants. The challenge of creating effective policies gets closer to the heart of the problem with traditional DLP: most solutions were designed to enforce static rules. Static rules can't easily account for anomalous behavior or activities that only appear risky when understood in context.

Survey participants also listed a lack of visibility into abnormal communication patterns among their top concerns, with 51% mentioning this inhibitor. Lacking visibility is an especially serious problem for organizations that experience more than 20 misdirected email incidents per year. 60% of respondents in that group ranked a lack of visibility into communication patterns among their top challenges.

Security executives were particularly likely to cite inadequate remediation workflows (63%) and time-consuming incident investigations (63%) among their top inhibitors.

Only 14% of participants reported that small team size was an inhibitor. And most (74%) said that inadequate budget was not an issue. Excessive false positive rates (mentioned by only 26%) were also not a top concern.

Which are the top factors limiting the success of your organization's efforts to prevent data loss through misdirected email? Select four.

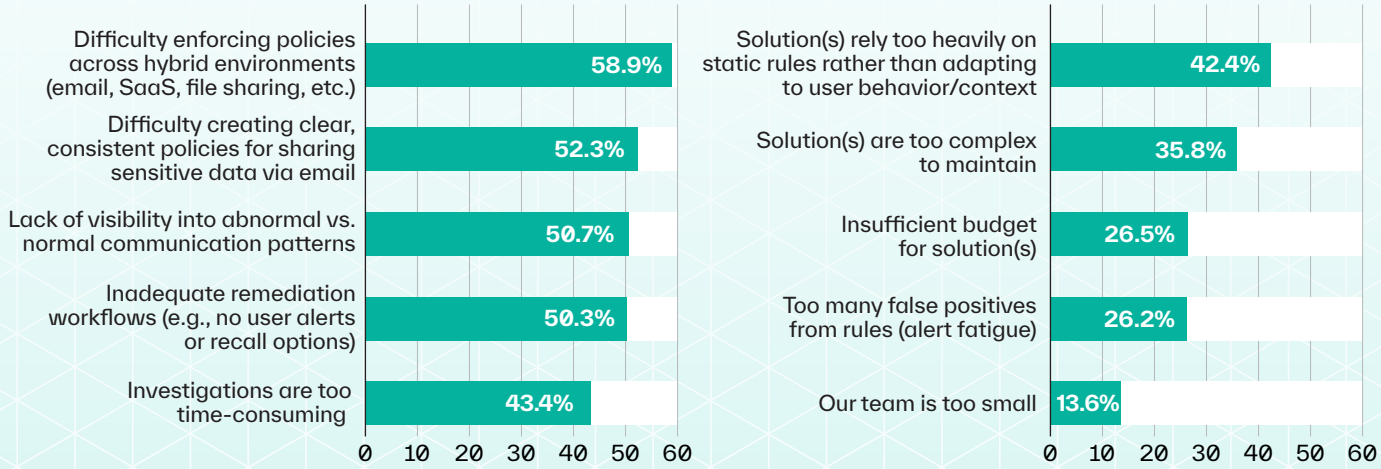

False Positives Drain Resources

Even though false positives weren't listed among the top concerns of survey participants, they're still taking up a great deal of security teams' time.

Most organizations (82%) are dedicating more than six labor hours each week to managing false positive alerts from the solution they use to prevent data loss through misdirected email. Just over one-third (36%) are spending in excess of eleven hours per week managing false positives, and one in ten (10%) need at least a half-time employee (someone working 20+ hours per week) just to keep up with their false positives.

This means that the average organization is spending more than 400 hours each year on managing false positives from its DLP or email security solution.

Earlier (see p. 12), we learned that excessive false positives aren't among the top obstacles to successful misdirected email prevention (perhaps because inaccurate detections have the potential to be even more damaging—and costly—than false negatives). But this is still a major problem for time- and resource-constrained teams.

400 The average organization spends **more than 400 hours** per year managing false positive alerts from its DLP or email security solution.

How much time does your organization spend on managing false positive alerts from the solution it uses to prevent data loss through misdirected email?

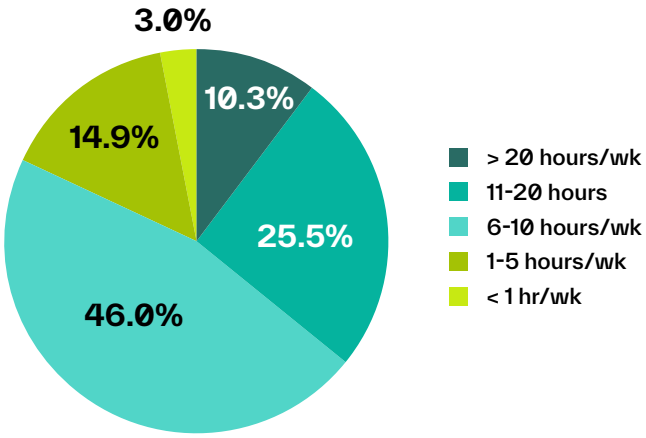

A New Formula for Success

Simply put, teams are looking for a solution that works. They're looking for capabilities that will enable effective, preventative blocking of risky communications. And they're looking for technologies that will get the job done without placing an additional burden on security teams.

Respondents recognize that behavioral AI will play a leading role in effective misdirected email prevention in the future.

Improving Misdirected Email Protection

We asked survey participants which capabilities would be most beneficial to them in a misdirected email prevention solution. The top response (mentioned by 69%) was the automated blocking of emails containing sensitive data sent to unintended recipients. Leveraging behavioral AI to identify anomalous data sharing or communication patterns (mentioned by 57%) was also highly ranked.

These capabilities go hand-in-hand. In order to automatically block employees from emailing sensitive data to unintended recipients, a solution must intelligently and accurately identify such scenarios consistently. This is exactly what behavioral AI enables.

Survey participants also ranked real-time contextual reminders for end users highly (this was mentioned by 51% overall). Contextual reminders are seen as particularly desirable among security executives (mentioned by 53%) and respondents in larger organizations (mentioned by 58% of those in organizations with 10,000 to 25,000 employees).

Lower priorities among survey participants include rapid deployment (mentioned by only 13%), seamless integration with other tools (mentioned by 20%), and granular policy controls (mentioned by 26%).

Ease of deployment and maintenance is less important to respondents than the solution's overall effectiveness.

Which of the following capabilities would have the greatest beneficial impact on a solution for reducing the risk of misdirected emails? Select three.

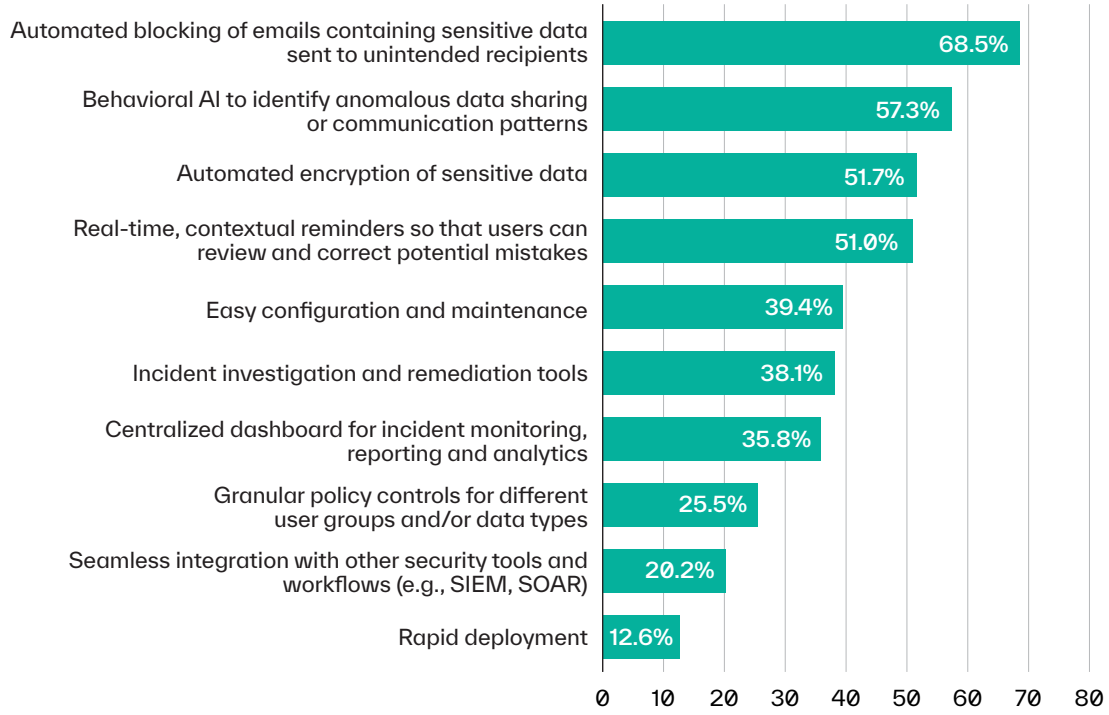

The Status Quo Is Not Working

Today's IT and security professionals are largely dissatisfied with the status quo when it comes to misdirected email protection.

A large majority (89%) agree that their current solution for preventing data loss through misdirected email requires substantial effort to configure and maintain. This is true even though 61% of respondents (see p. 15) don't include ease of configuration and maintenance among their top priorities for risk reduction. This administrative burden might not be among stakeholders' biggest problems, but it's still a problem.

Most respondents (77%) also agree that their organization's current solution for preventing misdirected email generates excessive false positive alerts. Even though these false positive alerts aren't their biggest obstacle to successfully preventing misdirected email (see p. 12), they're still diverting security teams' time and effort away from higher-value tasks.

About two-thirds (61%) of survey participants agree that their current solution isn't effective for preventing data loss through email. Given the prevalence of malicious data exfiltration via email—email plays a role in 61% of data breaches—it is unsurprising that current solutions aren't working.

A significant majority (70%) of respondents agree that the effort required to manage and operate the solution they use to prevent data loss through misdirected email isn't worth it.

Describe your agreement with each statement:

Our current solution for preventing data loss through misdirected email requires substantial effort to configure and/or maintain.

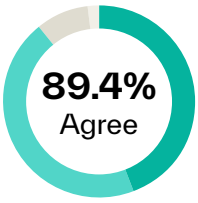

Our current solution for preventing data loss through misdirected email generates excessive false positive alerts.

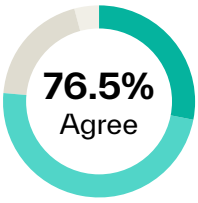

Our current solution is ineffective when it comes to preventing data loss via email.

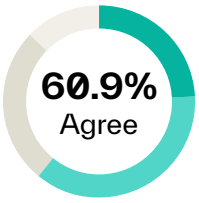

The effort required to manage and operate the solution we use to prevent data loss through misdirected email outweighs the value it provides.

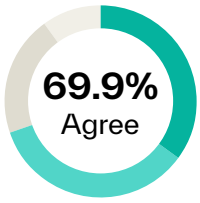

Strongly agree | Somewhat agree | Somewhat disagree | Strongly disagree

It's Time for Something New

Security stakeholders know the status quo isn't working, and they're ready to make a change. They are also prepared to invest in solutions that can achieve it. A full 81% of the participants in our survey said that their organization's budget for data loss prevention solutions is expected to increase over the next two years. For more than one-third of respondents (34%), that increase will be substantial (greater than 25%).

Nearly all survey participants (92%) expect that their organization will either increase or maintain their current level of spending.

The takeaway is clear: the market is ready for a new solution that can address IT and security teams' most pressing pain points—without adding operational overhead or more noise.

Which option best describes how you expect your organization's budget for data loss prevention to change over the next two years?

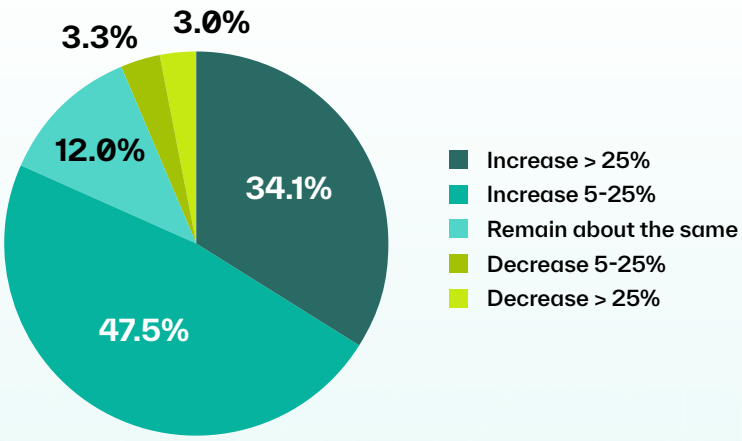

Behavioral AI Is the Answer

Most of these professionals are already aware that technology exists to solve the misdirected email problem. Survey participants broadly agree that behavioral AI is the key to effective misdirected email prevention.

The vast majority of respondents (94%) agree that applying behavioral AI and analytics in a solution to prevent data loss through misdirected email would significantly decrease the frequency of false positive alerts. All (100%) security executives and architects/engineers who took part in our survey agree with the above statement.

Nearly all (97%) survey participants agree that applying behavioral AI would significantly increase a solution's ability to detect truly risky data movements before they take place. This percentage increases to 100% among security executives, architects, engineers, and SecOps practitioners—arguably, the people best informed about AI's current capabilities and potential for improving operational workflows.

A similar percentage (96%) of survey participants agree that applying behavioral AI would significantly reduce the burden their team faces in defining rules and managing the solution. Again, this percentage increases to 100% among the architects and engineers responsible for overseeing this work.

The feedback is clear: security programs don't need more dashboards or complex rule sets. They need smart automation that solves the misdirected email problem before damage is done.

Describe your agreement with each statement. "Applying behavioral analytics and AI in a solution used to prevent data loss through misdirected email..."

... would significantly decrease the frequency of false positive alerts.

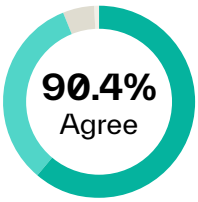

... would significantly increase its ability to detect truly risky data movements before they take place.

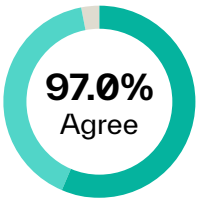

... would significantly reduce the burden our team faces in defining rules/managing the solution.

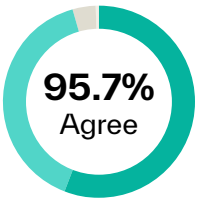

Strongly agree | Somewhat agree | Somewhat disagree | Strongly disagree

What Is Behavioral AI?

Behavioral AI is a type of artificial intelligence focused on modeling, analyzing, and understanding human (or entity) behavior to detect anomalies, automate threat response, and enhance security outcomes. It establishes individualized baselines for every user and account across the enterprise environment by analyzing thousands of factors such as communication patterns, login locations, and additional contextual signals. Behavioral AI continuously learns from normal behaviors so that it can detect and block abnormal, high-risk actions in real time.

- Because behavioral AI can reliably identify unfamiliar recipients, it can **flag high-risk emails before they are sent.**
- If a user attempts to send information outside a trusted domain, behavioral AI can **issue a real-time warning.**
- Behavioral AI solutions can support delayed sending windows to **enable last-minute recall** if a mistake is detected.
- Because it can identify and target only genuinely unusual behavior, behavioral AI **reduces alert fatigue for security teams.**

Abnormal Misdirected Email Prevention

Abnormal has developed a new behavioral AI-powered approach to misdirected email prevention, one built to stop accidental data leaks before they cause damage, without adding to the operational burden that today's busy security teams face.

Abnormal Misdirected Email Features:

AI Detection

AI-powered behavioral context and email metadata analysis identifies misaddressed communications without rules or policy tuning, instantly blocking misdirected messages and routing them to Microsoft Quarantine.

End-User Remediation

End users receive clear email notifications with release and delete options, eliminating SOC involvement.

Zero-Overhead Resolution

Direct decision routing to end users eliminates traditional SOC intervention bottlenecks.

Clear Audit Trail

All detections are logged in the Outbound Log with detailed context, sender-recipient metadata, and resolution status.

Conclusion

Misdirected email is a high-impact threat—one that's costing enterprises billions in remediation expenses and regulatory penalties—but its seriousness is often underestimated. Because these incidents stem from human error rather than malicious activity, it's all too easy for decision-makers to dismiss the risk, especially since an effective solution has long seemed elusive.

Traditional email security and DLP solutions lack visibility into the context surrounding communications, but a nuanced understanding of this context is what's needed to detect misdirected email before it causes harm. Without this kind of visibility, security tools will deliver an overwhelming flood of false positive alerts. Yet they'll still fail to prevent sensitive data from being sent to the wrong recipient.

This facet of email security and data protection is ripe for disruption. By harnessing the capabilities of behavioral AI to accurately detect high-risk emails before they're sent, a next-generation solution can finally solve the misdirected email problem. And it can do so while lightening the load on hardworking security teams and reinforcing safer communication habits across the organization.

Before the advent of behavioral AI, misdirected email was difficult for rule-based systems to identify. After all, the risk comes from non-malicious everyday communications simply sent to the wrong person. But today's advanced AI models are powerful enough to identify deviations from individual end users' typical behavior, and proactively quarantine messages before any damage is done. This can save hours of remediation time and hundreds of thousands of dollars in fines, and can also prevent devastating losses of customer trust that may never be repaired.

Abnormal's behavioral AI models human behavior with great precision. The ABX engine baselines normal activity and autonomously detects and prevents behavioral anomalies before they become breaches. Abnormal created Misdirected Email Prevention to set a new standard in email security, advancing the use of AI to safeguard organizations from human error.

About Abnormal AI

Abnormal AI is the leading AI-native human behavior security platform, leveraging machine learning to stop sophisticated inbound attacks and detect compromised accounts across email and connected applications. The anomaly detection engine leverages identity and context to understand human behavior and analyze the risk of every cloud email event—detecting and stopping sophisticated, socially-engineered attacks that target the human vulnerability.

You can deploy Abnormal in minutes with an API integration for Microsoft 365 or Google Workspace and experience the full value of the platform instantly. Additional protection is available for Slack, Workday, ServiceNow, Zoom, and multiple other cloud applications. Abnormal is currently trusted by more than 3,200 organizations, including over 25% of the Fortune 500, as it continues to redefine how cybersecurity works in the age of AI.

Discover how Abnormal closes the gaps in outbound email security.

See a Demo ›

Discover Your ROI ›

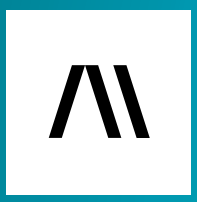

ABNORMAL.AI