


Abnormal

RESEARCH REPORT

The State of Email Security in an AI-Powered World

How Security Leaders are Responding to the Generative AI Threat



Executive Summary

98%

of security stakeholders are at least somewhat concerned about the cybersecurity **risks posed by ChatGPT, Google Bard, WormGPT**, and similar tools.

80%

of security stakeholders have confirmed that their organizations have **already received AI-generated email attacks** or strongly suspect that this is the case.

94%

of survey participants generally agree that **AI will have a major impact** on their organization's security strategy within the next two years.

46%

of respondents **lack confidence in traditional email security** solutions to detect and block AI-generated email attacks.

92%

of respondents agree that **“good” AI is valuable and necessary** for countering the risks posed by “bad” AI.

The rise of generative AI has been nothing short of meteoric. Tools like ChatGPT and Google Bard are becoming increasingly popular among people who wish to use them for legitimate purposes—inspired by the efficiency they create.

However, cybercriminals are also embracing this technology. Generative AI is likely behind the increases in both the volume and sophistication of email attacks that organizations have experienced in the past few months, and it's still early days. Attackers are already using it to create more realistic phishing emails without spelling or grammar errors, send highly persuasive business email compromise attacks, and write polymorphic malware. Unfortunately, the threat will only grow more severe in the future—and security leaders are taking note.

To understand how leaders are thinking about this threat, particularly in the email space, we surveyed 300 cybersecurity stakeholders from organizations of all sizes across multiple industries. Participants held positions at varying levels of seniority, but nearly one-quarter (24%) were CIOs, CISOs, or other senior leaders. Almost three-fourths (73%) of the survey respondents were at the manager level or above.

In this report, we dive into the survey results to uncover their concerns regarding generative AI, how they expect email security to evolve to respond to the threat, and how they're preparing their organizations for the future.

Table of Contents

The Lightning Fast Rise of Generative AI	4
Security Stakeholder Perceptions of Generative AI	10
Perceived Impact of Generative AI on the Email Threat Landscape	16
Plans for the Future: The Need for Good AI to Fight Bad AI	21
Conclusion	29
About Abnormal	30



The Lightning-Fast Rise of Generative AI

Over the past few months, we've witnessed the start of a technological revolution. [ChatGPT](#), OpenAI's large language model-based chatbot, has been making headlines ever since a free version of the tool was introduced to the public in November 2022. Almost instantly, ChatGPT shattered adoption records, attracting more than one million registered users within five days of its launch.

A plethora of similar AI tools, from [Google Bard](#) to [Microsoft Bing Chat](#), were rapidly introduced as competitors angled to capture part of this burgeoning market. In fact, [Gartner](#) predicts that by 2025, generative AI will be creating a full ten percent of all new data that's being produced.

A dark blue circle containing the white text "AI".

AI



From Zero to Hero: How Generative AI Is Used

The rapid adoption of these tools is no surprise. Business leaders and employees alike are eager to harness this technology's vast potential to create efficiencies—from drafting and summarizing documents to personalizing marketing campaigns, forecasting customer needs, and much more. People are even using it in their personal lives to budget, plan their travel adventures, and improve their written communications.


The number of use cases for generative AI will only grow as it's increasingly widely implemented. Unfortunately, though, at least some of these use cases are malicious. And because these tools are focused on generating text (as opposed to images or video), a particularly popular use case is creating email-based attacks—which were already seeing success even before generative AI was introduced.

Our survey participants are understandably worried that today's email-based threats will be amplified through generative AI, and they're especially concerned about the threats they see most often, like phishing and malware.

66%

of advanced email threats are **credential phishing** attacks.

Greatest Risks of Generative AI on Email Threats (Scale 1-5)



When asked about the greatest risks of generative AI, the responses were fairly tightly clustered across the board—indicating that all of the above-listed threats are a worry for security leaders.

And unfortunately, they should be worried. Bad actors are already taking advantage of this technology to create and disseminate large volumes of seemingly realistic email attacks.



Credential phishing, which accounted for two-thirds of all advanced email attacks even before ChatGPT's release, just got easier. "Spray and pray" credential phishing campaigns—i.e., those that are generic and target as many victims as possible in hopes that someone will be fooled or distracted enough to click on the malicious link—can now be created more quickly, allowing attackers to reach far more targets in less time.

What's even more concerning is that generative AI can make targeted socially-engineered attacks like business email compromise (BEC) and vendor email compromise (VEC) far more convincing. This renders them more likely to persuade targets to transfer funds or share sensitive information, and it also means they'll be far more difficult to detect by the end user.

Email was the most common first step in data breaches even before generative AI came onto the scene, but this technology has clear potential to increase the volume, sophistication, and resulting effectiveness of email-based attacks. Email attacks have historically posed major risks to organizations of all sizes across all industries—and now, things are only going to get worse.

Generative AI-Driven Attacks Are Already Underway

Security leaders have good reason to fear AI-generated attacks. While it is impossible to determine with certainty whether an attack was created by a human or AI, an assumption can be made that the continued increase in attacks in 2023 has been at least partially a result of the rise of generative AI.

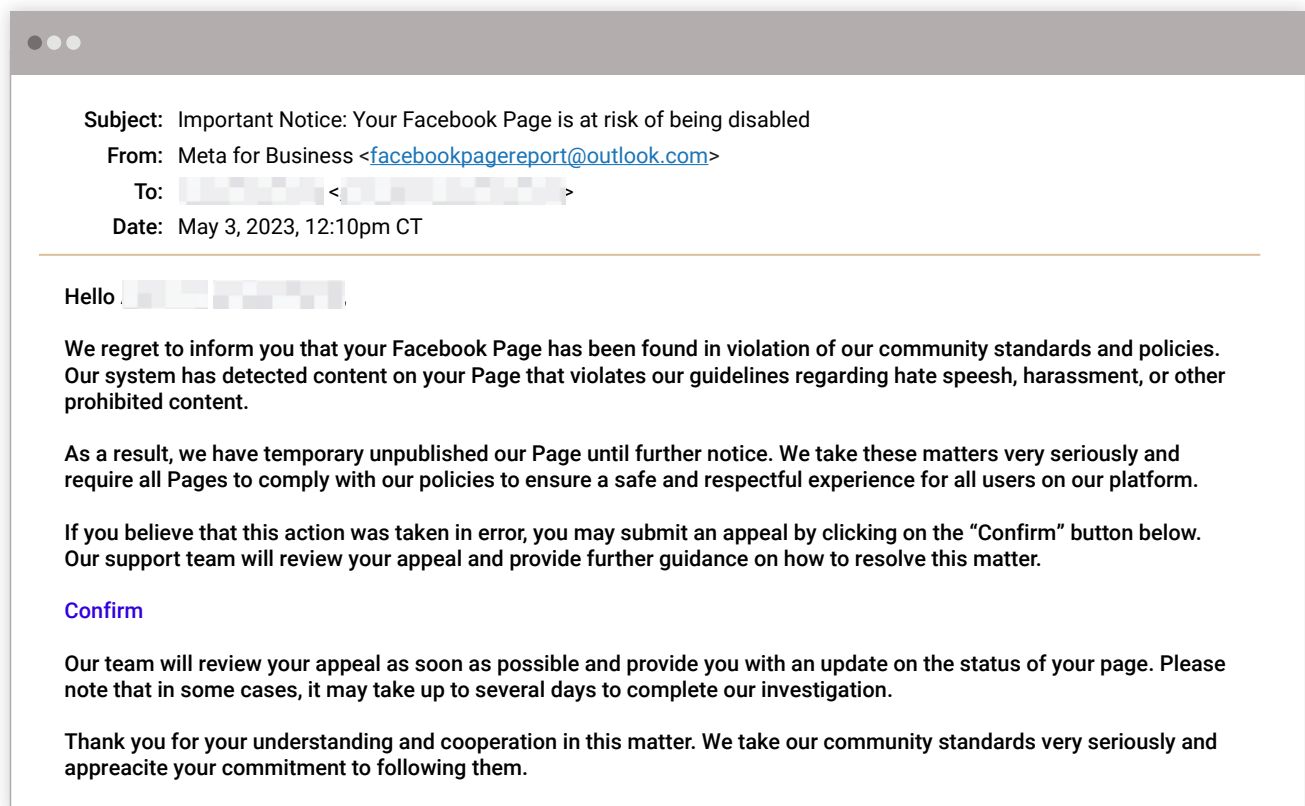
Part of the challenge in confidently determining causation lies in accurately detecting AI-generated attacks. The only solutions capable of accurately detecting AI-generated content are those that themselves leverage AI, using multiple LLMs and a suite of AI detection tools to make comparisons. Even so, these advanced solutions can only tell how likely it is that an email attack was generated by AI—not make an absolute determination.

Abnormal Security has recently developed CheckGPT, a tool able to make high-confidence determinations about whether or not a malicious email message is likely to have been created by AI. Here are some examples of likely AI-generated email attacks that we've detected within the past few months.

Credential Phishing via Brand Impersonation

This email sent by “Meta for Business” states that the recipient’s Facebook Page was found to be in violation of community standards and that the Page was unpublished. To resolve the issue, the recipient is informed they must file an appeal by clicking on an included link.

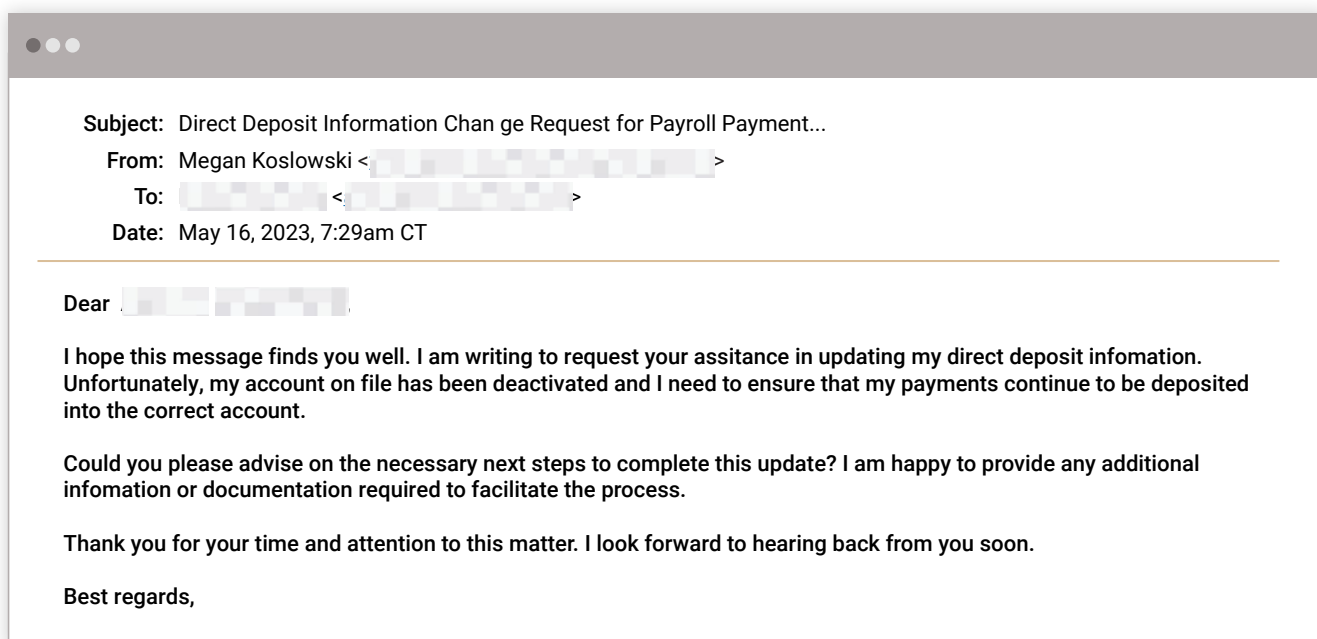
Of course, the link actually leads to a phishing site. If the user were to input their credentials, they’d immediately give attackers access to their Facebook profile and corresponding Page.



Payroll Diversion via Executive Impersonation

Generative AI is also being used in more targeted business email compromise attacks. These social engineering attacks rely on text-only emails and omit traditional indicators of compromise, like malicious links or bad sender domains, to bypass existing security protocols. They instead rely on name recognition and urgency to encourage the recipient to complete the request.

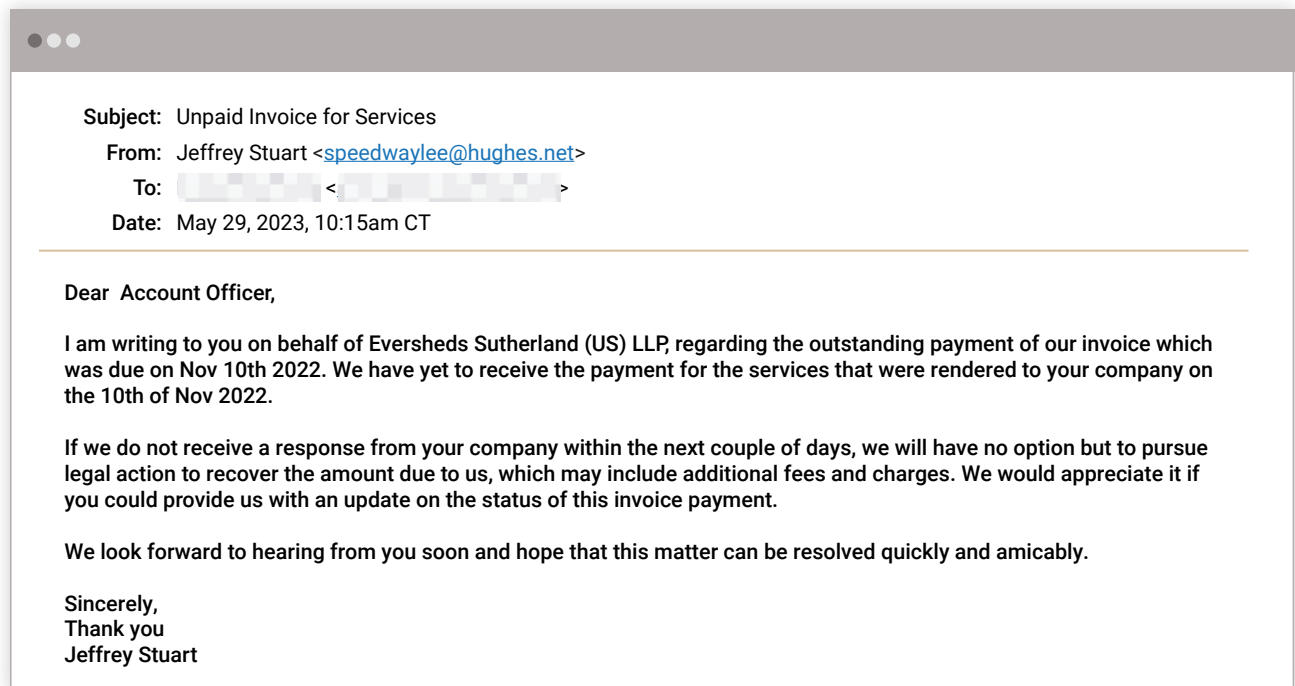
In the below attack, an executive's account is impersonated in a request to the payroll department to update the banking details on file for the next direct deposit. The email is free of grammatical errors or typos and is written very professionally—eliminating the telltale signs of attack.



Invoice Fraud via Vendor Impersonation

Attackers are also using generative AI in their vendor fraud attacks, where they impersonate business partners to request payment for fake invoices. Such attacks are among the most successful because they exploit the trust that's already present in relationships between vendors and customers. And because discussions with vendors often involve issues surrounding invoicing and payments, it's hard for end users to detect attacks that mimic these conversations.

This vendor fraud email involves the impersonation of an attorney who appears to be requesting payment for an outstanding invoice and is, again, free of spelling or grammatical errors. The impersonated attorney is also from a real-life law firm—a detail that gives the email an even greater sense of legitimacy and makes it more likely to deceive its victim.



It's clear that generative AI is already on the scene—and likely here to stay.
So where does that leave us on the security front?



Security Stakeholder Perceptions of Generative AI

With generative AI making headlines for the past several months and AI-generated attacks already targeting organizations, we expected to see security leaders express concerns about this technology's potential to exacerbate email-based attacks. That's exactly what our survey found.

Security stakeholders are currently concerned about generative AI's impact—and they're even more worried about what the future may bring as this technology further evolves.

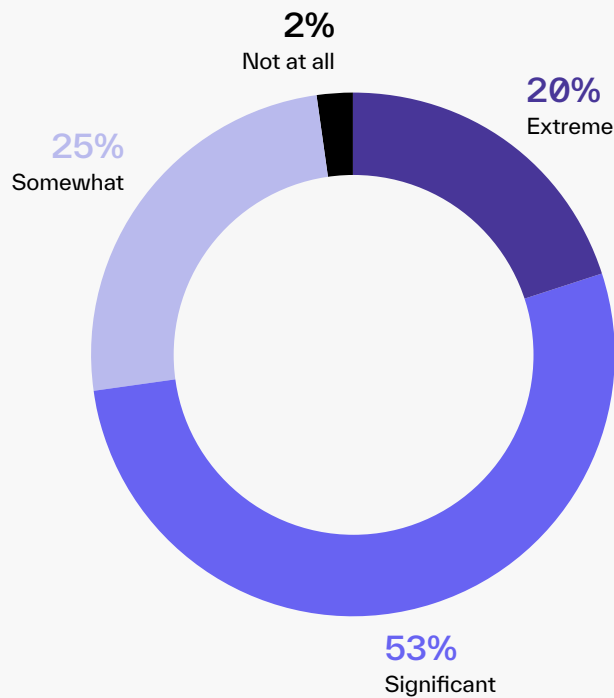


Concerns About Generative AI Are Rising

When asked how concerned stakeholders within their organizations are about the security risks currently posed by generative AI, most respondents voiced significant levels of concern.

Nearly all respondents (98%) are at least somewhat concerned and almost three in four (73%) said their concerns were “significant” or “extreme.” A mere two percent of respondents were not concerned at all.

Concern Regarding the Security Risks of Generative AI Today



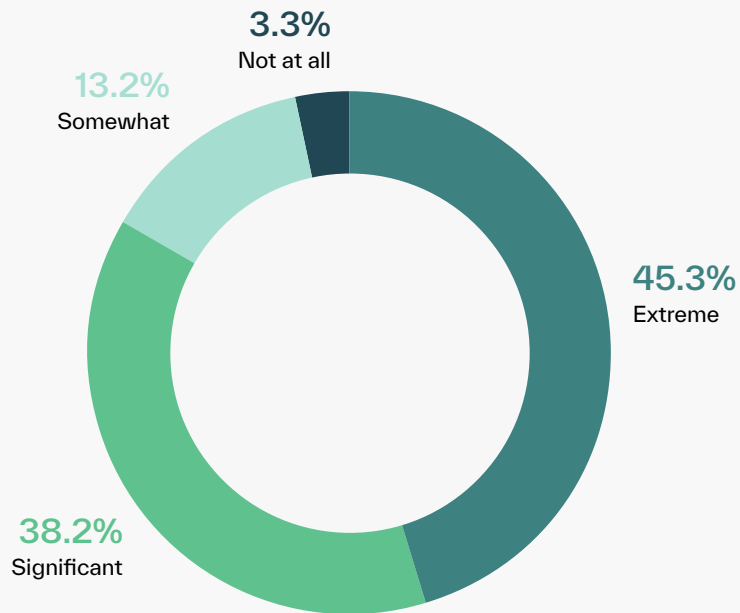
73%

of security stakeholders are "extremely" or "significantly" concerned about generative AI risks today.

These worries did not vary substantially across organizational size. For those companies with fewer than 2,500 employees, 71% of respondents were significantly or extremely concerned. This number rose to 74% for midsize organizations of 2,500-9,999 employees and 72% for the largest enterprises of more than 10,000 employees.

As concerned as they are about the risks that generative AI poses in the present day, these survey respondents are even more worried about the future.

Concern Regarding the Security Risks of Generative AI in One Year



97%

of security stakeholders are at least somewhat concerned about the **risks of generative AI in the next twelve months.**

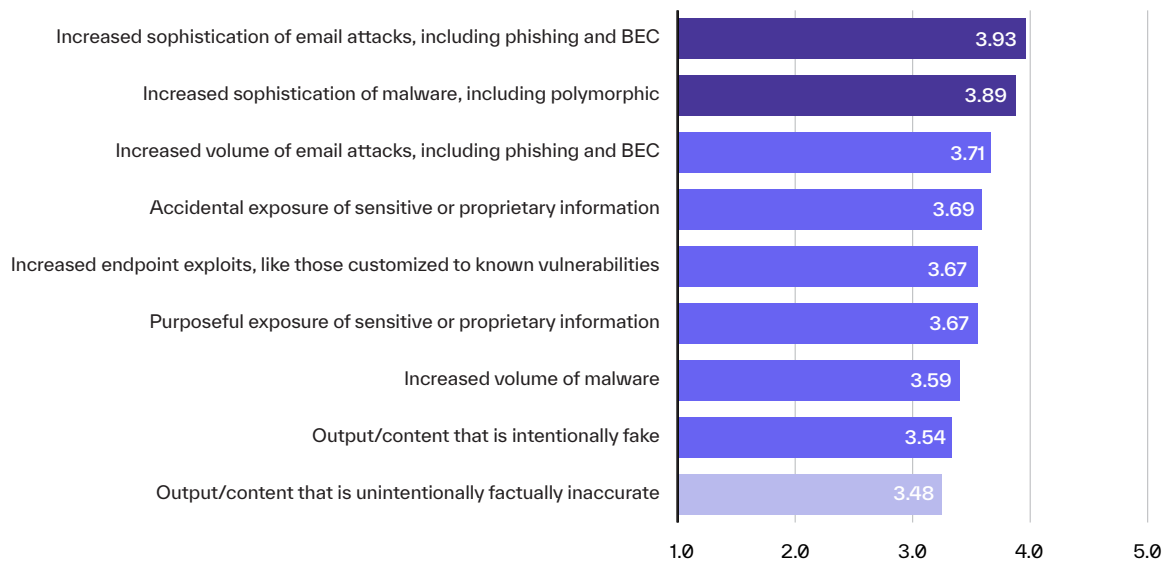
When it comes to the risks that generative AI may pose a year from now, security leaders are even more worried: 83% of respondents rated their level of concern as “extreme” or “significant” when asked about these risks.

These worries are certainly justified, as we know that attackers are already taking advantage of generative AI to move faster. Generative AI enables everyone to work more efficiently, and as cybercriminals see these attacks pay off, they’ll continue to launch more of them. What makes these tools particularly dangerous is that they can also open the door for less-savvy criminals to create technically-complex attacks, without the need for specialized knowledge.

Increasing Sophistication Is Among Top Concerns

Security leaders and practitioners recognize that generative AI is an emerging threat, making it impossible to know exactly how it will be used. That said, they're most worried about the sophistication of attacks, rather than volume.

Concern for Specific Risks Associated with Generative AI (Scale 1-5)



Routine Security Awareness Training

has long been a standard ingredient in enterprise cybersecurity toolkits.

The survey reveals levels of concern that are generally high across all areas. It's notable that even the lowest-rated risk of "output/content that is unintentionally factually inaccurate" is still above the midpoint of the spectrum, meaning that respondents are still more (rather than less) concerned about it.

These results are unsurprising. After all, every one of these issues represents a real-world threat that is likely to become even more pressing in the coming months.



If generative AI renders security awareness training ineffective, what other measures will organizations need to adopt to supplement it or replace it?

Among security leaders' main areas of concern, two stand head-and-shoulders above the rest: the increased sophistication of the email attacks that generative AI will make possible and the increased sophistication of the malware that it can be used to create.

Both are valid worries. Using tools like ChatGPT or its infamous cousin [WormGPT](#), would-be cybercriminals with minimal technical skills can easily create both sophisticated email attacks, as seen in the examples shared earlier in this report, and malicious software. While ChatGPT itself has built-in content filters intended to prevent it from being used to create malicious emails and malware, [these can be easily bypassed](#).

[Security researchers](#) have also found that ChatGPT can be used to create an endless number of new versions of a snippet of code—making it highly capable of creating polymorphic malware that is evasive and difficult to detect. What's more, the ChatGPT API can be leveraged within malware itself, meaning it can deliver modules to perform different actions as needed. This results in malware that does not exhibit malicious behavior or suspicious logic.

Another leading concern—one with a great deal of merit—is that the volume of email attacks is likely to increase with the adoption of generative AI. This will put greater pressure on end users: will they be able to accurately identify malicious emails in growing numbers? It also has the potential to add to the already significant workload that security teams must manage. If attack volumes increase, and these attacks become more convincing, it's all but certain that the burden upon analysts will grow, too.


Further, generative AI tools stand to make security awareness training much less effective. Traditionally, these trainings have instructed end users to look out for the classic signs of an email attack, like misspellings, grammar or syntax issues, and typos. With generative AI, it's possible to engineer an email attack in which the grammar and spelling are absolutely flawless—and where the writing style perfectly mimics a brand voice or the way a coworker whose account has been compromised would typically communicate.

Routine security awareness training has long been a standard ingredient in enterprise cybersecurity toolkits. If generative AI renders security awareness training ineffective, what other measures will organizations need to adopt to supplement or replace it? The need for more robust inbound email security—and stronger built-in protections against account takeover attacks—is clear.

Concerns Abound for Multiple Email Attack Use Cases

Security stakeholders have far-ranging concerns about a wide variety of AI-powered email attack tactics, but they're most nervous about the fact that that generative AI will help attackers craft highly specific and personalized email attacks based on publicly available information.

Concern by Attack Use Case (Scale 1-5)



Here again, we see survey participants recognizing and expressing legitimate concerns. Responses are fairly tightly clustered, indicating that all of the attack use cases asked about are creating apprehension. However, security professionals are most worried about the ability to create highly specific attacks, identify additional vulnerabilities, and automate tasks.

Generative AI can do all of these things, and it can do them quickly—without requiring bad actors to have technical skills or expertise. AI tools can almost instantly surface and summarize relevant details about a prospective target (including their social media posts, public records, and other online content) to help attackers personalize the malicious messages. They can also scan code for known vulnerabilities and dramatically accelerate the process of sending large volumes of email attacks.

Less concerning for respondents, but still worrisome, is the ability of attackers to use generative AI to analyze content from compromised accounts to identify relevant details from conversation histories. This enables threat actors to copy a victim's tone or discover intellectual property to steal. Generative AI can also be used to effortlessly translate phishing emails or BEC attacks into multiple languages, with results that are grammatically correct and highly realistic. The mistakes that attackers typically make when writing in a language that's not their own will no longer be a telltale sign of malicious email content.

The question now remains: **what can be done?**






Perceived Impact of Generative AI on the Email Threat Landscape

One of the many challenges in defending against these threats is the fact that AI allows attackers to create net-new attacks at scale—rendering solutions that look for traditional indicators of compromise useless against this new threat. **So how are security leaders thinking about upleveling their email security tools in this new AI-powered environment?**

Solutions in Today's Inbound Email Security Stacks

We started out by asking survey participants what tools and technologies their organizations had in place to protect employees against email threats.

Email Security Tools Currently in Use



Reliance on the traditional secure email gateway has declined over the last year, from 59% to 53% as more organizations rely on the functionality available from the native email security provider.

The majority of survey respondents (77%) are relying on the native features and functions that their email provider—Microsoft 365 or Google Workspace—makes available to them. This marks a 15% increase from our findings in [last year's survey](#) and is likely to be expected given the increased focus on budget as macroeconomic conditions shifted this year.

Traditional secure email gateway (SEG) usage has declined slightly, from 59% last year to 53% this year, as growing numbers of security stakeholders become aware of its limitations in detecting socially-engineered attacks. **Many are instead relying on native security features provided by Microsoft as the provider continues to innovate and now offers comparable capabilities.** In addition, more than 40% of all organizations are using an integrated cloud email security solution (ICES) in addition to either their native security functionality or their secure email gateway.

ICES adoption is highest among small and midsize organizations, suggesting that these companies may be more agile and less burdened by technical debt. Among organizations with fewer than 10,000 employees, 41% have implemented an ICES solution, while only 38% of enterprises with more than 10,000 employees have done so.

As AI-generated email attacks continue to grow in prevalence, their rise may compel greater numbers of companies of all sizes to put an ICES solution in place. These solutions—unlike secure email gateways—integrate directly with cloud email providers, so they have immediate access to thousands of signals needed to detect socially-engineered attacks. This extensive visibility into the email environment makes it possible for ICES solutions to understand the content of an email message within the context of an organization to better detect suspicious activity and filter out malicious emails—no matter who created them.

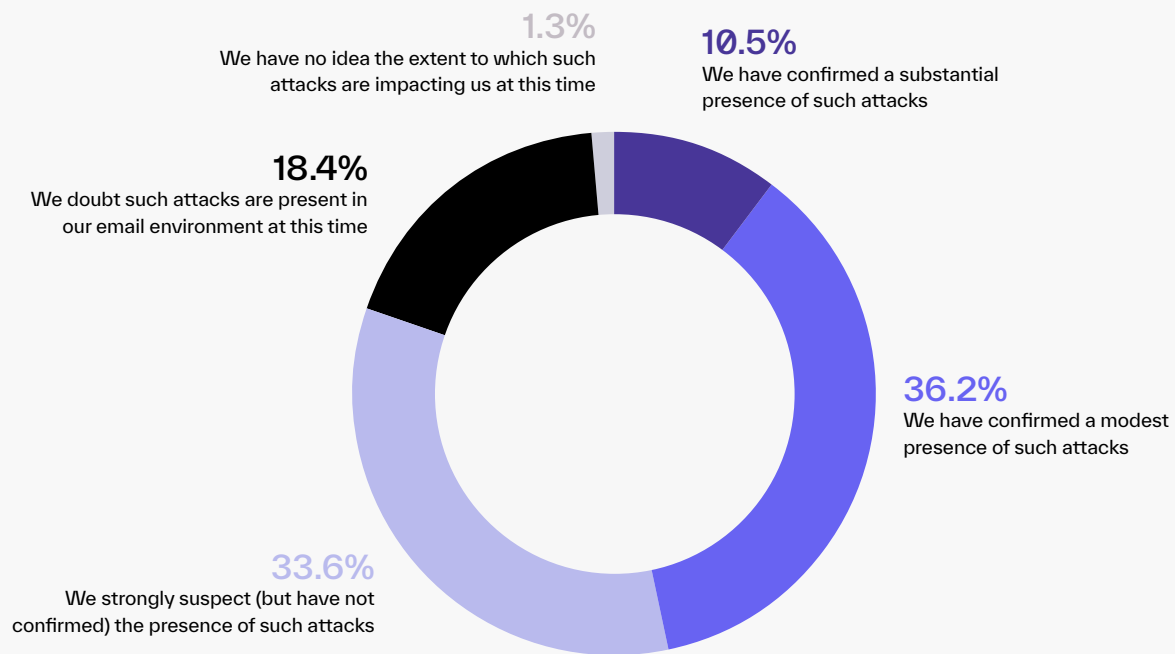


Security Professionals Already Observing AI-Generated Email Attacks

Understanding the email security architecture for organizations is increasingly important in the age of AI, as a clear majority of survey participants believe (or have confirmed) that their organizations are seeing AI-generated email attacks.

While it's hard to be entirely confident that these responses are accurate, given that even the generative AI models themselves struggle to identify AI-created content consistently and accurately, it is a clear indicator that this fear is no longer hypothetical. Ultimately, what's more important than identifying malicious AI-generated content is the ability to stop it—regardless of how it was created.

Level of Exposure to AI-Generated Email Attacks



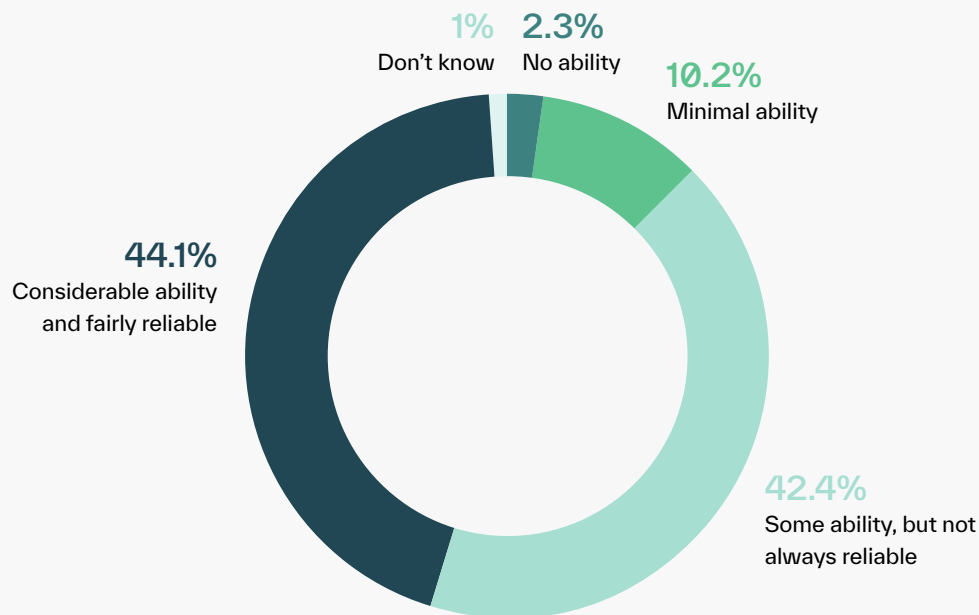
A clear majority of respondents (80%) either have confirmed that their organizations have received AI-generated email attacks or strongly suspect that this is the case. Only a minority of survey participants (18%) believe that they are not being targeted by these types of attacks, with an additional 1.3% unsure of the extent to which they have received AI-generated attacks.

Both responses, though, beg an important question: how can survey participants tell whether email content is generated by AI or by humans? Without access to sophisticated analytical models that can determine the probability that a particular email message was generated by AI, most business and security stakeholders are making only educated guesses about the impact that generative AI is truly having on their organization.

How Email Security Solutions Distinguish AI-Generated Attacks

We asked survey participants whether their current email security solution could tell AI-generated email attacks apart from ones that were written and launched by humans. It's likely that their responses display a certain degree of overconfidence, given that very few email security solutions that are currently available actually have this capability.

Email Security Solutions' Ability to Distinguish AI-Generated Email Attacks



57%

of organizations with an integrated cloud email security solution believe that it can detect AI-generated attacks.

Still, as many as 86% of survey participants believe that the email security solution that's currently in place within their organization has at least some ability to distinguish AI-generated email attacks. And a full 44% of respondents believe their organization's email security solution has the ability and is fairly reliable. In contrast, only 13% of respondents believe their organization's email security solution has little or no ability to distinguish AI-generated email attacks from those produced by humans—a percentage that should likely be much higher.

It is worth noting, however, that respondents from organizations that have implemented an ICES solution are significantly more likely to have confidence in the technology's ability to distinguish AI-generated email attacks from those not created using AI. In fact, 57% of survey participants in organizations with ICES have confidence in the solution's ability to detect AI-generated attacks, whereas only 32% of those using a SEG believe that the SEG can accomplish this.

98%


of survey participants are at least somewhat confident that AI-powered security solutions are needed to detect and block AI-generated threats.

Stakeholders Are (Perhaps Overly) Confident in Technology and Users

Recognizing that email security tools are not the only way to mitigate the impact of AI-generated attacks, we asked survey participants about their confidence levels in a variety of other tools and processes.

In general, respondents have a moderate to high degree of confidence in the email security solutions in use in their environments. There's little doubt that at least some of this confidence is misplaced, as email-borne threats remain a major problem for organizations both large and small. Even before the massive rise of AI, business email compromise alone resulted in losses of **\$2.7 billion** last year—a number that is likely to continue to increase.

Security Professionals' Confidence in...



Can You Identify an AI-Generated Email Attack?

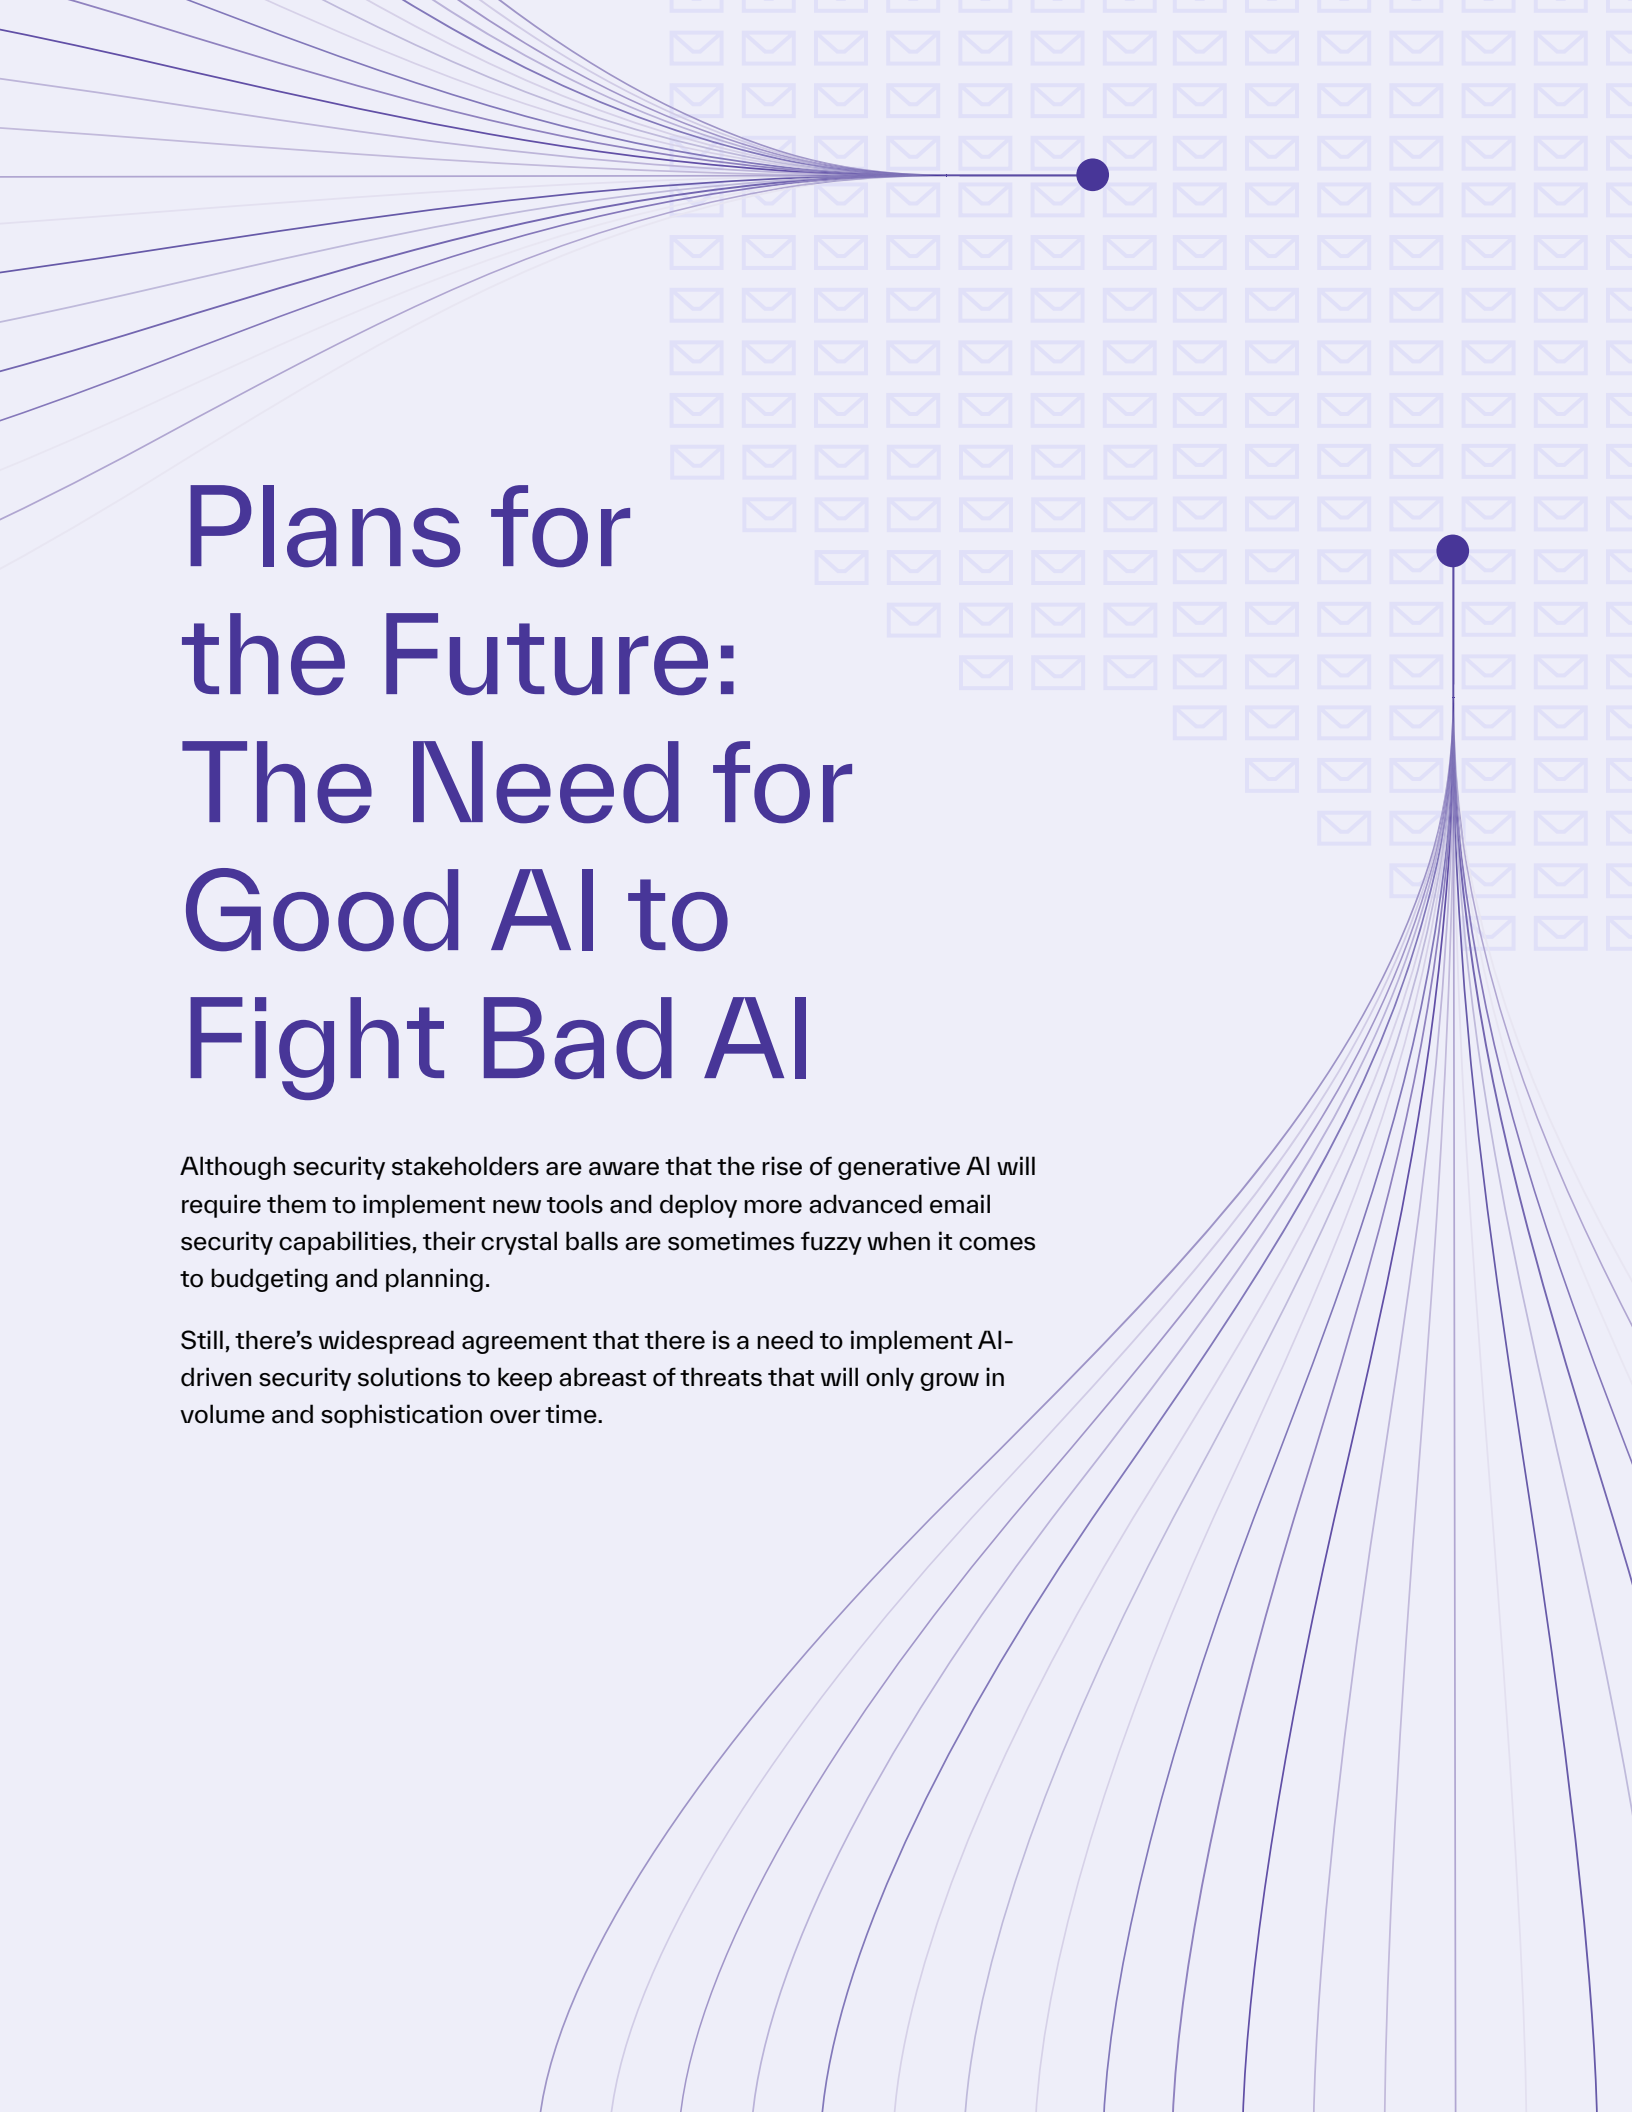
Test your ability to distinguish between attacks written by humans and ones generated by ChatGPT.

Take Quiz →

That said, respondents tended to agree with the fact that AI-driven security solutions are likely better able to detect and block AI-generated email attacks, with a full 61% significantly confident in this technology's ability to give defenders an edge.

It's interesting to note that 50% of survey participants have significant or extreme faith in the ability of end users to identify and avoid these attacks, and another 48% believe that government or industry regulations could solve the problem. In both cases, this likely remains to be seen.





Plans for the Future: The Need for Good AI to Fight Bad AI

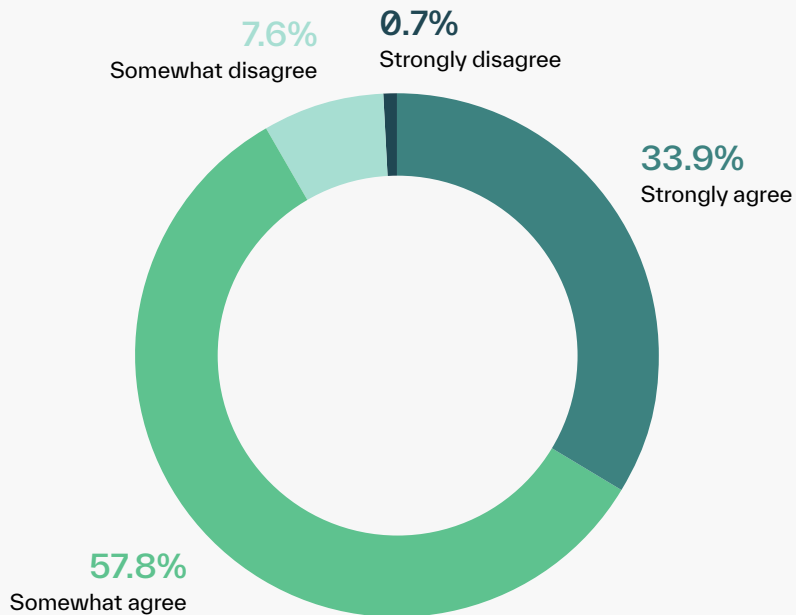
Although security stakeholders are aware that the rise of generative AI will require them to implement new tools and deploy more advanced email security capabilities, their crystal balls are sometimes fuzzy when it comes to budgeting and planning.

Still, there's widespread agreement that there is a need to implement AI-driven security solutions to keep abreast of threats that will only grow in volume and sophistication over time.

AI-Driven Security Is Needed for Effective Defense

A majority of survey participants do believe that AI-driven security solutions will be needed to defend against AI-generated threats, and most are confident that AI will have a major impact on their organizations' security strategies within the next few years.

AI-Driven Security Is a Must-Have



A clear majority of survey participants (92%) see value in using AI-driven security solutions to defend against today's AI-generated email threats. And less than one percent of respondents strongly disagree with that viewpoint.


Senior executives are particularly likely to agree, with a full 96% of the CIOs, CISOs, or higher-level technology leaders participating in this survey saying they strongly or somewhat agree.

AI Is Already Transforming Enterprise Security Strategies

To better understand how the rise of generative AI is changing how leaders and practitioners plan for the future of their cybersecurity programs, we asked survey participants whether AI will have a major influence on their organization's cybersecurity strategy within the next two years.

Given the results of the previous question, it is no surprise that there is widespread agreement that AI will impact their organizational strategy. More than 94% of survey participants say that AI will have a major impact on their cybersecurity strategy over the next two years.

AI Will Have a Major Impact on Cybersecurity Plans for the Organization in the Next Two Years



97%


of CISOs and executives believe that AI will have a major influence on their cybersecurity strategy moving forward.

Both those who work with these technologies in a hands-on fashion and those who are responsible for setting long-term strategic goals understand the revolutionary impact that these technologies are having on the field of IT security. Over 97% of executive leaders see the value of AI in cybersecurity in the next two years, and 100% of email and messaging managers believe that they will be moving to more AI-focused solutions. This is also true for frontline information security practitioners like security analysts and incident responders who were in 97% agreement with the statement.

Talent and Skills Shortages Holding Security Programs Back

So what is keeping organizations from making the shift to better defense today? To understand this, we asked survey participants to rate a number of potential barriers to building effective defenses against AI-generated email attacks. We wanted to understand which inhibitors were blocking progress the most, as well as which obstacles were seen as less problematic.

Greatest Inhibitors to Defending Against AI-Generated Email Attacks



3.4 million

cybersecurity positions remain unfilled around the world.

Broadly speaking, respondents believe that a shortage of security professionals is the greatest barrier to successfully defending against today's advanced AI-generated email threats. Nearly as important is a general lack of knowledge and skills pertaining to AI.

Survey participants also recognize that there's a need to leverage AI to counter AI-driven threats. Without taking this step, defenders will be relegated to playing an endless game of catch-up with attackers—leaving them to take reactive approaches that are doomed to fail. That said, perhaps the reason why they are using these measures insufficiently is due to the lack of personnel required to manage them.

Interestingly, insufficient budget was the least-mentioned issue. It seems that when it comes to countering AI-generated email attacks, or sophisticated email attacks in general, organizations are willing to invest.

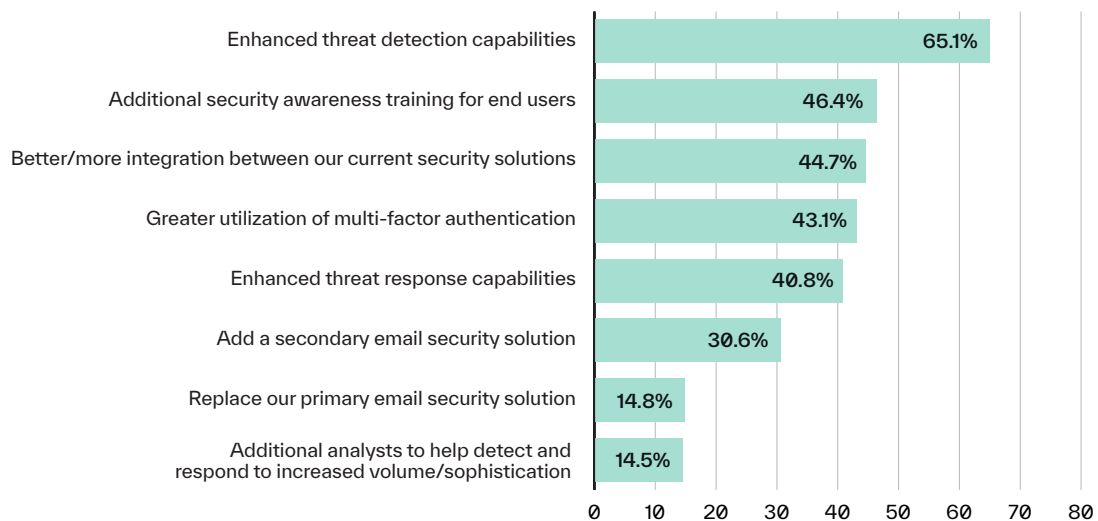


Shifting Short-Term Priorities for Defending Against AI-Generated Email Attacks

While it's clear that participants in our survey understand the transformative impact that generative AI is having on the email threat landscape and want to start improving their defenses, it's less clear whether they have formulated strategies for dealing with the risks in the short term.

When asked what their specific plans are for addressing the issue, many survey participants reported that they'll simply continue as they've been operating. That being said, some did report intentions of finding new tools or upgrading existing ones to better harness the power of AI.

Top Priorities for Defending Against AI-Generated Email Attacks




31%

of security professionals are looking for an additional email security solution to help block AI-generated attacks.

The largest group of respondents (65.1%) intend to invest in enhancing their organization's threat detection capabilities. Nearly half of survey participants (46%) are also planning to invest in additional security awareness training for end users. While the latter is certainly an essential element of cybersecurity, it is somewhat of an odd decision for this particular threat—given how difficult it is for humans (trained or not) to identify highly sophisticated, authentic-seeming emails.

In 45% of cases, security professionals plan on more tightly integrating the solutions in their security stack, which could help detect lateral attacks or better indicate compromised accounts. In addition, leaders intend to utilize multi-factor authentication—potentially helpful in thwarting credential phishing attacks—and enhance threat response capabilities to stop attacks more quickly.





Relatively few respondents are planning to replace their organization's primary email security solution. This is understandable, given the costs associated with implementing a new solution and the desire to make the most of what has already been invested. However, a larger number of survey participants expressed willingness to invest in a second email security solution—likely in an effort to provide additional layers of defense to detect and block more attacks.

In addition, a limited number of respondents plan to add additional security analysts to their teams. This may well be due to the shortage of personnel mentioned in the previous question. In today's world, most organizations may instead be seeking to supplement their lean teams with technology solutions that can empower them to work more efficiently and effectively.



The Shift to AI-Native Email Security

Over two decades ago, pioneers like Salesforce attempted to challenge the status quo of on-premises software by introducing a novel concept: **cloud-native** software. These companies had a vision of delivering value by harnessing the full power of the cloud. They built teams and products to realize that vision, with every line of code written specifically for the cloud. As organizations moved more and more of their operations to the cloud, these pioneers delivered a disproportionate amount of value to their customers—allowing them to leapfrog the competition.

Today, we're seeing the same phenomenon take place with artificial intelligence. Companies are increasingly seeking out AI-native technologies, which were created with AI from the very beginning, and leveraging their power to reimagine operations. As AI transforms the way the world uses software and how bad actors launch attacks, it can also be used to improve how organizations protect themselves.

AI-native cybersecurity companies are those that place AI at the core of their protective capabilities. As the business world and bad actors alike move in the AI-powered era, these are the companies poised to lead the revolution in an effort to use good AI to fight bad AI.



The Need for Good AI to Combat Bad AI

For better or worse, generative AI and the attacks it enables are here to stay. Security leaders must be prepared, both now and in the coming years, to protect their organizations against this emerging and evolving threat.

Secure Email Gateways vs. Integrated Cloud Email Security Solutions

Unfortunately, the traditional tools developed for email security were designed for a different era, when organizations maintained on-premises servers and followed a rules-based approach to finding and blocking malicious emails. The secure email gateway still reroutes emails for analysis and examines them for known indicators of compromise like suspicious URLs and malicious attachments. But with the rise of both social engineering tactics and the use of generative AI, it is now nearly impossible for these tools to stop the full spectrum of attacks.

In contrast, integrated cloud email security solutions connect directly to cloud email providers via API—ensuring access to thousands of signals and data from across the email environment to detect suspicious activity. The tools leverage machine learning and behavioral AI to baseline known-good behavior and identify anomalies, relying on identity modeling, behavioral and relationship graphs, and deep content analysis to highlight emails that appear suspicious. This AI-based technology can take into consideration a broad array of factors—such as internal and cross-organizational relationships and geolocation, device usage, and login patterns—to detect malicious behavior even when there are no traditional indicators of compromise present.

Not only can AI make it possible for an email security solution to understand what's normal for each individual sender and recipient using sophisticated technology like natural language processing, natural language understanding, and image recognition, but it can also detect atypical logins or communication patterns. This fundamentally different approach is focused on identifying deviations from a behavioral baseline (signaling potential fraud or malicious intent), which allows it to detect and block attacks of all kinds.

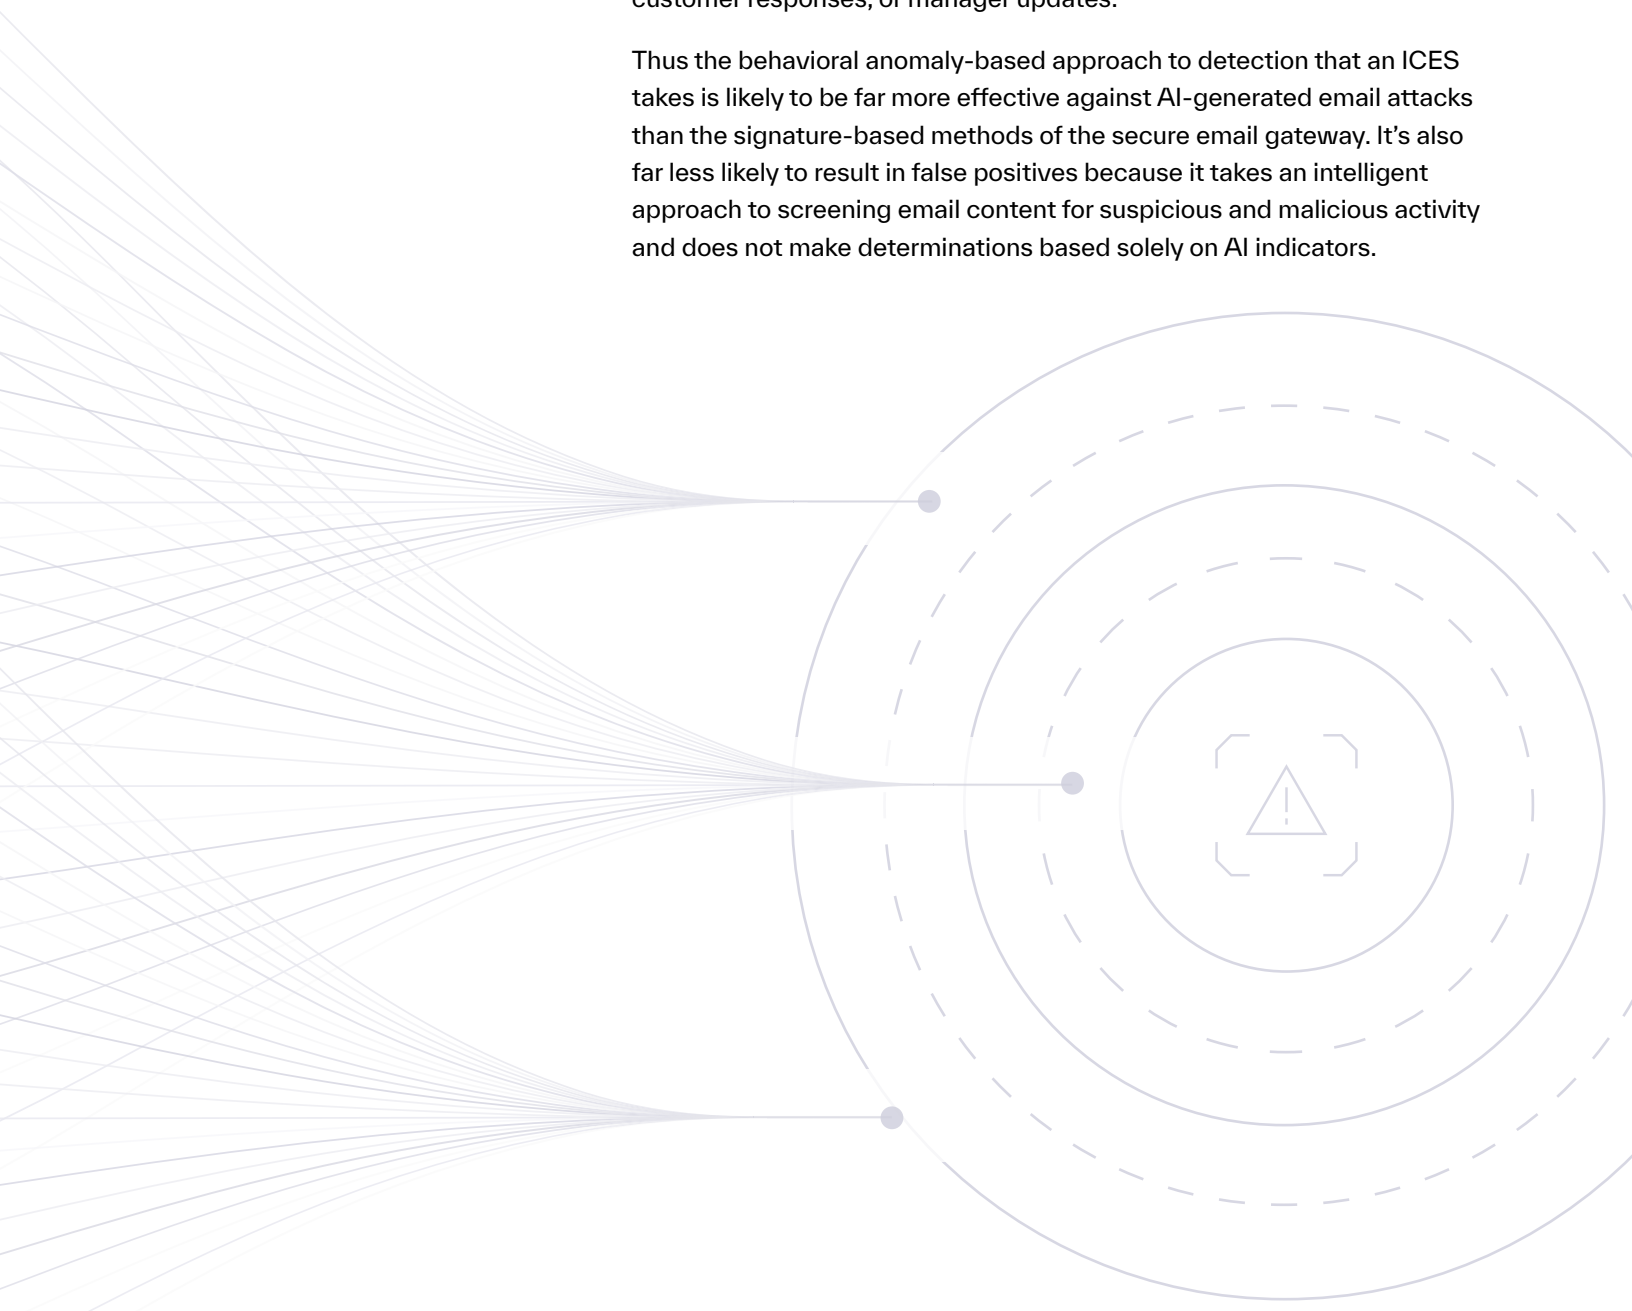


Using AI to Detect AI

This distinction between the architecture of the two solutions is important to understanding how they initially detect AI-generated messages. That said, whether or not email security solutions are able to detect AI-generated email content is far less important than whether they're able to detect and block malicious emails effectively—no matter how they were created.

Even if an email security solution can consistently and accurately identify AI-generated emails, simply blocking them all isn't the answer. This would interfere with legitimate business communications as employees rely on generative AI to help them craft email messages for marketing content, customer responses, or manager updates.

Thus the behavioral anomaly-based approach to detection that an ICES takes is likely to be far more effective against AI-generated email attacks than the signature-based methods of the secure email gateway. It's also far less likely to result in false positives because it takes an intelligent approach to screening email content for suspicious and malicious activity and does not make determinations based solely on AI indicators.



Conclusion

Attackers have long adapted their tactics to evade defenses, moving from the Nigerian Prince scams of the past to more targeted business email compromise and from basic viruses to polymorphic malware.

Unfortunately, the same premise holds true for generative AI. Because tools like ChatGPT and Google Bard promise to deliver enormous business value, their use cannot simply be forbidden. Nor can enterprises just block every email message that appears to have been created by AI.

Security stakeholders recognize this as a concern and are (rightly) worried about how generative AI will increase the volume and sophistication of email attacks. When it comes to mitigating these growing risks, there's still much work to be done. Today's security programs remain overly reliant on traditional security solutions, and stakeholders are overconfident in the ability of both tools and end users to recognize AI-powered attacks.

This means a new approach is needed: one that leverages the power of AI to stop email attacks, regardless of whether they were generated by AI models or created by humans. The good news is that security leaders recognize this fact and understand the potential that good AI has to counter bad AI. This is an important step in the right direction.

The next step is widespread implementation of AI-native email security solutions. When implemented correctly, these advanced solutions can prevent the full spectrum of email attacks with extremely high efficacy, detect and remediate compromised accounts, and streamline operations to increase productivity—providing better experiences for end users and security teams alike. Plus, by using defensive AI to detect attacks, they can ensure that your organization stays protected from today's most pertinent threat: malicious AI.

Survey Details

This survey of 300 security professionals was conducted via a third-party in August 2023. It included organizations of all sizes across multiple industries.

Participants held positions at varying levels of seniority, but 24% were executives while 73% were at the manager level or above.

Abnormal

Abnormal Security provides the leading behavioral AI-based email security platform that leverages machine learning to stop sophisticated inbound email attacks and dangerous email platform attacks that evade traditional solutions. The anomaly detection engine leverages identity and context to analyze the risk of every cloud email event, preventing inbound email attacks, detecting compromised accounts, and remediating emails and messages in milliseconds—all while providing visibility into configuration drifts across your environment.

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, with additional protection available for Slack, Teams, and Zoom.

More information is available at abnormalsecurity.com

**Interested in Detecting and Preventing
AI-Generated Email Attacks?**

Request a Demo →

See Your ROI →