


Sublime Email Threat Research Report 2026



Introduction

Welcome to the 2026 Sublime Email Threat Research Report, a retrospective of the threats we've observed across our customer base in the past year.

This report looks back at 2025, examining how email-based threats have evolved throughout the year across organizations of all sizes. Our analysis of detection data across customer environments highlights several trends that security practitioners should be aware of to maintain effective defenses.

Contents

2 Audience Takeaways

3 Executive Summary

5 Core Topics

Continued evolution of BEC and social engineering: Lowdown on hijacking

Threat uniqueness and customization: Eyes on impersonation

Evasion techniques by attack type: Questioning QR codes

Evasion stacking: TTPs all the way down

Getting hot: Opportunistic LOTS, ICS phishing, email bombs

25 Final Thoughts

Conclusion

Recommendations

Methodology

Audience Takeaways

This report is designed for security leaders and practitioners responsible for email security in their organizations. It provides:

- 1 Actionable intelligence on emerging email threat tactics, techniques, and procedures (TTPs)
- 2 Industry-specific threat analysis to help tailor defensive strategies
- 3 Technical insights into evasion techniques used by threat actors
- 4 Data-driven recommendations for security control improvements
- 5 Insights into emerging attack vectors to inform security roadmaps

Executive Summary

Through 2025, we observed email threats shift from volume-based attacks to precision-engineered campaigns that exploit human psychology and organizational trust. This is in keeping with a broader shift in threat actor tradecraft over the last several years toward dynamic behaviors intended to evade static detection and interdiction across all phases of an intrusion. Our analysis of millions of messages from the past year has revealed trends in the methods attackers are employing. We will dive deep into five core threat topics, but here is a high-level summary of the past year in email attacks:




FIGURE 1 Attack type distribution by %

BEC + fraud

Business email compromise (BEC) and fraud were the dominant threat category, **accounting for nearly one in three confirmed threats (32.0% of total threat volume)**. These attacks increasingly leveraged thread hijacking and fake threads (28.1% of observed BEC) to deceive or manipulate human targets.

Credential phishing

Credential phishing was the **second largest threat (23.6% of total threat volume)**. We observed attackers increasingly using malicious calendar invites and QR codes as an out-of-band pivot. In the case of QR codes, the email itself was often intentionally sparse and mainly served to move targets off the inbox and onto a personal mobile device, where corporate controls and monitoring are often weaker.

Evasion stacking

Evasion stacking is becoming the norm, **with 34.7% of all observed attacks employing two or more evasion techniques, and new emerging techniques including a ~100x increase in ICS phishing**. This indicates that traditional layered security approaches that rely on blocking individual attack vectors are insufficient. Organizations need behavioral analysis that can detect social engineering patterns regardless of the technical delivery method.

GenAI attacks

GenAI attacks show trends of both growth and evolution. By analyzing sample sets of malicious messages for specific indications of likely AI-generated content, **we saw a leap from 4.21% in Q1 to 19.29% in Q4**. This significant leap is likely due to both increased adoption of attacker AI and the rapid iteration/evolution of AI attacks which makes certain signals flare up and fade away. In either case, it is a clear indication that organizations need to be prepared for the speed and scale of AI attacks and their evolution.

Our data revealed a clear trend: **attackers are optimizing for trust, not just volume**. As detection capabilities improve, threat actors focus on smaller, highly targeted campaigns that exploit psychological vulnerabilities, such as thread hijacking and VIP impersonation, to manipulate targets into lowering their guard.

STANDOUT STATS

28.1%

of BEC attacks featured
thread hijacking

~10%

of attacks had specific
signals of potential
AI use

~100x

increase in ICS phishing
from Q1 to Q4

12%

of all attacks abused
trusted sites

>20%


of Q4 callback phish
were delivered via
calendar invite

10,000+

emails have been caught
in a single email bomb

Core Topics

- 1 Continued evolution of BEC and social engineering: Lowdown on hijacking
- 2 Threat uniqueness and customization: Eyes on impersonation
- 3 Evasion techniques by attack type: Questioning QR codes
- 4 Evasion stacking: TTPs all the way down
- 5 Getting hot: Opportunistic LOTS, ICS phishing, email bombs



Continued evolution of BEC and social engineering: Lowdown on hijacking

As seen in the chart above, BEC/fraud attacks represented ~32% of emails in 2025 – the largest threat category by volume. Unlike credential phishing's straightforward approach, BEC is defined as an attack that manipulates targets through impersonation, urgency, and exploitation of trust. Here's a breakdown of the top techniques used in BEC and fraud attacks:




FIGURE 2 BEC/fraud attack distribution by %

The data reveals an evolution: **thread hijacking and fake threads now comprise 28.1% of all BEC attacks**, making it the leading BEC technique. Attackers insert themselves into legitimate email threads, often through compromised accounts, creating the appearance of continuing an authentic business discussion. This makes detection challenging: the sender domain may be unfamiliar, but the context and conversation history appear genuine.

This faked context reduces friction for other tactics, like:

Spear phishing

Most phishing attacks appear as cold outreach, initial emails meant to start a conversation from an unknown sender. A more advanced shift in operations is for attackers to insert themselves into existing conversations or generate a thread that appears legitimate within the organization.

VIP impersonation

Using generative AI, it has become comparatively trivial to target a key position within the organization. In targeting a VP of Finance, for example, the attacker could spoof a realistic thread with the names and addresses of executive leadership, suggesting that the target is being included in an ongoing conversation. If a finance officer sees days of fabricated previous discussions regarding sending money for a purchase order, there is inherent trust that the activity is legitimate.

Inside access

In a more sophisticated attack, threat actors already have access to the environment and can read and write from specific mailboxes. Using an existing thread, they can bring new targets into the email chain under the guise of continuing authentic business discussions.

Thread hijacking and fabrication make detection challenging for both humans and automated systems. The threads look legitimate and trustworthy, and many targets would shy away from directly questioning executive leadership to verify that an email chain is genuine.

The evolution of BEC from cold call to hot hijack

Tired: Traditional BEC

- Unsolicited email from "CEO" requesting wire transfer
- No conversation history or context
- Easily detected by recipient skepticism and security tools

Wired: Modern BEC

- Reply to existing email thread about ongoing project
- Appears to continue legitimate business discussion
- Often received without skepticism due to plausible, but fake context

BEC and fraud attacks by industry

While thread hijacking was the top tactic across most industries, advance fee fraud stood out far above the rest, accounting for **43.9% of attacks in Agriculture**. This may be attributed to large payments for equipment, feed, etc., prevalent within the Agriculture industry.




FIGURE 3 BEC category distribution by industry (% within each industry)

Threat uniqueness and customization: Eyes on impersonation

Modern attackers have begun shifting away from generic campaigns toward highly customized, organization-specific attacks. Analysis of our dataset reveals that **18.9% of BEC attacks showed significant impersonation elements**, suggesting attackers are either investing substantial effort in researching targets and crafting convincing personas or leveraging GenAI to convincingly accomplish the task.

While the concepts behind each are similar, the nuances of how attackers use them can hint at who the intended targets are and the attacker's sophistication.




FIGURE 4 Attacks featuring impersonation (% by type)

Brand impersonation is the most common and basic style of deceit. Emails reuse logos, styles, and layouts of well-known and trusted sites to trick a user into believing the content is actually hosted from that site. This may be a malicious file hosted from a site masquerading as Dropbox or a fake login page that is nearly identical to one used by Microsoft.

The top impersonated brands were
**SendGrid and the IRS, combining for nearly
10% of all brand impersonations.**

10%

VIP and executive impersonation attacks are a common and well-known form of targeted manipulation. Attackers will take on the role of a company's executive leadership and make requests of the target to send information or money. Instead of using proper channels, the attackers will spoof text message numbers or send from a personal email address with stories of how the VIP is stuck without access to their official email but has an urgent need for gift cards.




FIGURE 5 VIP/executive impersonation attack subcategories by %

Employee impersonation is similar to VIP impersonation in that it also mimics someone within the same organization. However, these tend to be more crafted and specialized, with the attacker knowing more intimate details about the organization's structure and its employees. Without the power dynamic of a VIP, their emails require believable and trusted context in order to succeed. This could be via a fake WorkDay page sent to a member of HR, or an attacker impersonating an IT technician with a malicious link related to a recent application update.





FIGURE 6 Employee impersonation attack subcategories by %

These attacks all require various amounts of reconnaissance on target organizations. Targeted attackers invest time in studying communication styles, organizational charts, technology adoption patterns, and timing of business processes to maximize attack effectiveness. However, GenAI is changing this.

GenAI attacks have allowed attackers to bypass much of the effort these attacks used to require. By providing LinkedIn profiles and company pages, an attacker can leverage AI to process this data instantly, build org charts, and create believable personas, backstories, and entire email threads. Social media accounts can be analyzed to determine a VIP's typical grammar and spelling, for example, to create legitimate-looking emails that mimic a CEO who does not capitalize words and uses new lines instead of punctuation.

Indications of growth in AI-generated attacks

While it's not possible for any technology to identify AI-generated content with 100% accuracy, we have identified a variety of key signals that are frequently found in AI-generated emails. Using a selection of these signals, we analyzed a subset of malicious emails to surface trends in the usage of offensive AI. Our analysis showed an upward trend of AI in attacks with a significant **increase up to 19.3% in Q4**.



The rapidly iterative nature of AI attacks can be seen in the way signals can suddenly appear and then quickly spike. Due to AI's ability to unlock attack evolution at speed and scale, it is important for organizations to use email security tools that can keep pace. Looking ahead, it seems clear that defensive agentic AI in email security – where AI agents use the latest tools and models to detect attacks and close gaps – will become table stakes in the near future in order to defend at attacker speed.



The trouble with classifying AI is that it only takes a prompt update to remove AI indicators from content. For example, students have learned to instruct LLMs to not use the em dash (—) in papers after that specific punctuation became a known indicator of AI for professors.

Impersonation tactics by industry

Executive impersonation

Industries with stronger executive power dynamics (Energy, Finance, etc.) and potential for highly-impactful, time-sensitive action requests (Utilities, Healthcare, etc.) show more instances of executive impersonation.


Employee impersonation

Real Estate and Hospitality are industries built upon peer trust, so attackers tend to leverage user-to-user trust rather than rely on a power dynamic.

Brand impersonation

Highly regulated industries – including the Public Sector – often have very restrictive controls on vendors and data sharing. Impersonating a random file sharing service will be less effective, as it would be easily detected as an exception. In contrast, educational institutions and professional services engage with a wide variety of outsiders and generally expect to receive a wide range of services in inbound emails.

FIGURE 8 Impersonation type distribution by industry



Evasion techniques by attack type: Questioning QR codes

Attackers employ sophisticated evasion techniques to bypass ever-hardening security controls. Our analysis identified distinct evasion patterns across different attack types.

Besides BEC and fraud attacks, our analysis also included credential phishing and callback phishing attacks. Together, these two types comprised 39.5% of all malicious emails. There were some common sub-categories for these two types of attacks that were also present in BEC and fraud attacks, most notably thread hijacking (10.1% of credential phishing attacks, and 6.0% of callback phishing attacks) and employee impersonation (9.1% of credential phishing attacks).




FIGURE 9 Credential phishing attack sub-categories breakdown




FIGURE 10 Callback phishing attack sub-categories breakdown

Spotlight on evasion techniques

CALENDAR INVITE ABUSE (ICS PHISHING)

8.5%

of callback phish were on calendars

8.5% of all callback phishing attacks used calendar invites to deliver malicious phone numbers and instructions, exploiting the fact that calendar systems often bypass traditional email security scanning.

LOTS (LIVING OFF TRUSTED SITES)

12%

of attacks abused services

12% of attacks leveraged legitimate cloud services such as Google Drive, OneDrive, and SharePoint, as well as emerging platforms like Jotform and LiveAgent Desk to evade infrastructure-based analysis. By hosting malicious content on trusted domains, attackers bypass numerous infrastructure reputation filters.

QR CODE EMBEDDING

1.4x

more likely a QR code was malicious

This technique uses QR codes embedded in emails, attachments, linked pages, cloud-hosted documents, and more to deliver a payload. When a QR code was observed, it was 1.4x more likely to be an attack than a message without a QR code.

COMMODITY PHISHING


57%

of callback phishing was old school

Attackers still heavily use the traditional techniques of "call me" in high volume. While new and sophisticated attacks were discovered during the year, we continued to see a high volume of well-known callback phishing.

QR code delivery methods

QR codes can be delivered a variety of ways, but we continue to see them primarily embedded directly within the body of an email. This makes it imperative that email security solutions recognize QR codes and interpret their destinations (through redirects and gates).



QR code attack growth from H1 to H2

Observed QR code attacks surged dramatically throughout 2025. This increase came with a substantial rise in QR codes delivered via stacked evasion.



Why attackers love QR codes

1. **Bypasses email gateway:** Traditional email security analyzes text and URLs, not images containing QR codes.
2. **Pivots to mobile:** Targets typically scan QR codes on personal phones, which operate outside corporate security controls.
3. **Stacked evasion:** LOTS-linked QR codes combine two evasion techniques (trusted domain + visual obfuscation).



Most dangerous combination: QR codes delivered via LOTS represent the most sophisticated attacks, requiring computer vision to detect QR codes in cloud-hosted documents, link analysis to identify suspicious cloud-hosted content, and mobile device security to block malicious URLs.

FIGURE 12 Email attack evasion techniques by Industry - Percentage breakdown



Evasion stacking: TTPs all the way down

Another trend observed is the rise of evasion stacking, attacks that employ multiple evasion techniques simultaneously to maximize the likelihood of bypassing security controls. Our analysis reveals that **22.7% of all observed attacks employed two or more evasion techniques** in the same campaign.


The evasion types we analyzed were:

- **File-based:** Embeds the payload in a file to avoid detection or delay execution.
- **Behavioral:** Evades detection by manipulating timing, patterns, or social context to appear legitimate.
- **Content-based:** Techniques that manipulate the visible body content to obscure intent or evade keyword/signature detection.
- **Encoding and localization:** Uses alternative languages, character sets, or encodings to bypass traditional filters.
- **Header-based:** Manipulations in email headers that obscure sender, recipient, or subject intent.
- **Infrastructure-based:** Tactics that manipulate sender infrastructure or the delivery path.

Evasion stacking varies by attack type

Adversaries will stack evasion methods to best suit their attack. In the case of extortion, we've observed the greatest variety of combinations, indicating that adversaries are employing a "kitchen sink" approach to these attacks. Other attack types show consistent combinations, indicating that adversaries have tuned their methodologies to have tried and true evasion stacks.


FIGURE 13 % of BEC attacks stacking Behavioral + Header-based evasions



BEC: Behavioral + Header-based

BEC attacks rely on social engineering (Behavioral) and impersonation (Header) to manipulate their targets into performing further actions.


FIGURE 14 % of callback phishing stacking Behavioral + File-based evasions



Callback phishing: Behavioral + File-based

To evade body scanning, callback phishing information is often embedded in attached documents (File), meaning the language in the body of the message must manipulate the target into launching the file (Behavioral).

FIGURE 15 % of malware/ransomware attacks stacking Behavioral + Content-based evasion



Malware/Ransomware: Behavioral + Content-based

In modern malware/ransomware attacks, while the file is the payload, it isn't usually the evasion. More often, adversaries will use manipulative language (Behavioral) and disguised links (Content) to deliver payloads from remote servers, not attachments.

Specific stacking

We then dug deeper on the above evasion types to determine which specific TTPs were most frequently combined in attacks.

LOTS + QR code

Attackers host a PDF on Google Drive containing a malicious QR code. The Google Drive URL passes reputation filters, the PDF passes malware scans, and the QR code redirects to a credential harvesting site only when scanned on a mobile device.

Brand Impersonation + thread hijacking

Attackers reply to an existing email thread while impersonating a known brand (e.g., DocuSign), combining the trust of conversation context with brand recognition.

SVG Files + LOTS

Malicious SVG files hosted on SharePoint or OneDrive, leveraging both the trusted domain and the ability to embed scripts within image files.

VIP impersonation + calendar invites

Executive impersonation combined with a calendar invite containing a callback number, exploiting both authority bias and the tendency of calendar systems to bypass email security.

The strategic implication is clear: traditional layered security approaches that rely on blocking individual attack vectors are insufficient. Organizations need behavioral and multi-stage analysis that can detect social engineering patterns regardless of the technical delivery method.

Getting hot: Opportunistic LOTS, ICS phishing, email bombs

1. Increased usage of lesser abused LOTS hosts and SaaS platforms

Traditional LOTS


LOTS attacks use sites as malicious middlemen, providing a legitimate face to illegitimate content. Hosted content can be in the form of a fake login page or simply a malicious file.

Because the initial content is being hosted from a well-known and trusted site, there is often minimal, if any, logging and alerting on files being accessed. We have also tracked many LOTS hosting providers that are used only to deliver malicious emails. While not hosting attacker's content on their websites, they leverage the trusted domain to send illegitimate links to targets. For example, an email sent via Google that has a link to a credential phishing payload on a non-Google domain.

Modern LOTS


While security teams are aware of abuse on major cloud platforms (e.g., Google Drive, OneDrive, Dropbox, etc.), we saw a noticeable shift toward newer and more specialized LOTS attacks that are often unknown to defenders. Analysis revealed that **32.8% of observed email attacks leveraged uncommon or emerging platforms** rather than well-known services.

FIGURE 16 LOTS attacks by site



Generally, the largest LOTS-abused sites are diligent in monitoring and removing these threats. The surge in new services by startups, alongside the rise of AI-assisted content tools, has created a long tail of small providers that lack the resources to effectively police user-generated content. This is evident in the rise in opportunistic LOTS attacks.

Online design platforms let attackers easily create and publish a page that looks identical to a Microsoft login page or a file download page.



These emerging platforms fall into a few broad categories:

- **Page and form builders:** Jotform, Typeform, Buildin.AI, Canva, and similar platforms allow attackers to create convincing credential harvesting forms, fake CAPTCHAs, and illegitimate file sharing pages.
- **Customer support platforms:** LiveAgent Desk, Zendesk, and Intercom can enable attackers to host phishing pages that mimic support ticket interfaces.
- **Document and file sharing:** WeTransfer, 1fichier, Uguu, and similar hosting sites allow for files to be temporarily stored and shared widely.
- **Collaborative workspaces:** Notion, Airtable, Coda, and other sites offer advanced team collaboration services, acting as one-stop-shops for email workflows, page building, and file sharing.


Trend implication:

The 32.8% share of uncommon platforms indicates attackers are actively diversifying beyond monitored services. Security teams cannot rely on static blocklists of "known bad LOTS providers" and must implement behavioral and agentic detections that identifies suspicious patterns regardless of hosting platform.

2. ICS phishing: Calendar invite abuse

One of the most significant developments in 2025 was the rapid adoption of calendar invitations for delivering email attacks. In Q4 alone, we observed that **over 20% of callback phishing attacks used calendar invites as one of its delivery mechanisms** in 2025.

FIGURE 17 Attacks featuring calendar invitation payloads



These attacks exploit the fact that meetings are often automatically added to a calendar when an invitation is received. Email clients often auto-add the message body and attachments to the calendar event, so malicious payloads exist in both the calendar and inbox.

Making this more dangerous, in cases where the phishing email is blocked, the meeting invite still reaches the target's calendar. Since calendar invites can bypass mail processing entirely, whether you have an SEG or API-based email security solution, you need special handling to effectively detect, prevent, and remediate these attacks.

What makes ICS phishing so effective?

- **Urgent:** Users who receive a sudden meeting reminder in the middle of their workday may be more likely to review the invite and open attachments to determine why they were invited.
- **Expected:** In many instances, the lure is a conference call invite, but with explicit instructions to call in via phone, connecting the target directly with the attacker.
- **Tried and true:** ICS phishing uses the same social engineering techniques that phishing attacks have been honing for decades, such as urgency, financial incentives, impersonations, etc., but are being delivered through a new vector that targets have not been trained to be suspicious of.

Beyond ICS phishing

While ICS phishing is trending upward, there are other techniques in the wild that can deliver this two-pronged attack (sometimes called cross-channel phishing). For example, assigned tasks in Google Docs can often send both an email and a push notification to the recipient's phone, guiding them to a malicious document.

These cross-channel attacks can create a false sense of security, as teams may assume that catching the email stops the whole attack, which allows the secondary payload slip through undetected. This can even encourage attackers to craft deliberately detectable emails knowing that it will be caught in service of hiding the secondary vector.



Sublime stops ICS phishing in the inbox and on the calendar.

[Read our ICS phishing blog](#) →

3. Email bombs: The DDoS of the inbox

Email bombing is the practice of overwhelming a target's inbox with hundreds or thousands of automated messages as an attack in itself or as a smokescreen for another attack. Email bombs emerged as a significant operational threat in the latter half of 2025.

Email bombs are effective due to being:

- **Fast enough** to overwhelm users (creating chaos and urgency)
- **Slow enough** to potentially evade rate-limiting thresholds
- **Long enough** to cause significant disruption
- **Short enough** to complete before manual intervention

10,000+

Largest email bomb analyzed

360.5

Average size of an email bomb

17.85

Average message per minute in an email bomb

Disable mailboxes

Attackers can use email bombs to disable a target mailbox or make it otherwise unusable due to the volume of messages.

Machine takeover

Attackers will initiate an email bomb, pose as "tech support" to "fix" the bomb, and then get the target to install a remote access tool.

Account takeovers

Attackers will initiate a password reset, use a bomb to hide the legitimate reset email, and then send a fake reset email that phishes credentials.

Where do the emails within an email bomb come from, and how do hundreds of unique emails from hundreds of known domains suddenly appear? In most cases, this is simply from websites that allow for online email signups to newsletters without using a CAPTCHA. This allows attackers to create tooling to automatically sign up a user to hundreds of forms at the moment of the attack.



Sublime builds patterns of email volume and behavior for each mailbox at an organization. Once an email bomb spike reaches a mailbox-specific threshold, all the messages in the bomb are grouped and auto-remediated.

[Learn more about email bomb prevention →](#)

Final Thoughts

- 1 Conclusion
- 2 Recommendations
- 3 Methodology



The 2025 email threat landscape is defined by two dynamics: the exploitation of trust and the rapid pace of change in attacker tactics.

Attackers are investing in context and, through context, exploiting trust at scale. These dynamics are mutually reinforcing a broader trend across intrusion tradecraft: namely, zeroing in on trust relationships as a critical point of potential failure within security postures. Just as many sophisticated actors target the trust invested in vendor relationships through supply chain compromises, phishing operators exploit trust in ongoing conversations through techniques like thread hijacking – all targeting the assumption that established relationships are inherently safe.

In practical terms, this operational evolution extends beyond just the tactic of thread hijacking. While the basic concept has dramatically diversified over the years, the data points to an accelerating trend of impersonation tailored to specific verticals' predominant workplace cultures. Attackers will target trust where it organically resides in business relationships, be it between executives and employees or between peers or with external vendors. AI tools allow them to do this with more believability, more organization-specific data, and more volume than ever before.

And to many users, what is more trusted than their calendars? As the organizing principle for daily workflows, the calendar represents a domain of implicit trust that attackers have rapidly moved to exploit, with calendar invite abuse growing from negligible in Q1 to almost 20% of callback phishing in Q4. This pattern extends beyond calendars: cross-channel techniques like ICS phishing with push notifications exploit the assumption that non-email platforms have already vetted the content they surface – yet another trust boundary that attackers are increasingly targeting.

The rate of additional tactical innovations is also significant. QR code attacks increased 282.7% between the first and second half of 2025. Email bombs emerged as both a disruption tactic and a smokescreen for credential theft. Evasion stacking, combining multiple techniques in a single attack, appeared in 34.7% of campaigns. Attackers are continuously testing new vectors, weaponizing the inherent trust of emerging LOTS platforms and the obscurity of under-monitored file types to probe for coverage gaps.

These dynamics compound into a singular challenge: attackers are eroding trust across more vectors than ever, while organizations still need communications trustworthy enough to function. Offloading the balance between warranted skepticism and operational paralysis onto users' judgment is an increasingly untenable position — that burden should rest on detection systems that are adaptive enough to counter both established and emerging tradecraft.

**Move as fast as attackers
with next-gen email security.**

Get a demo:

<https://sublime.security/demo/>



1 Combine deterministic detection with agentic AI for full-spectrum coverage

Traditional ML models remain essential for deterministic, low-latency detection of known patterns, but often struggle with low-frequency novel tactics and multi-step evasions without new training data. Organizations should deploy a hybrid detection architecture where behavioral rules, ML models, and AI agents complement each other.

ML models provide probabilistic signals and generalization to unknown permutations of bad behavior, rules provide deterministic enforcement and rapid iteration, and AI agents provide zero-day detection of never-before-seen techniques. The most effective systems allow rules to directly leverage ML outputs, and agents able to leverage rules, to improve precision without sacrificing explainability or response speed.

2 Operationalize response with AI-powered investigation, remediation, and adaptation

Operationalizing response requires delivering consistent end-to-end outcomes: detection quality, accurate prioritization, and transparency in decision making. Organizations should deploy AI agents that autonomously triage suspicious messages as volume rises, correlate threats, hunt for missed attacks, and write new detection coverage to close identified gaps in real time. The goal is to transform security operations from reactive investigation to continuous, high-fidelity protection that reduces risk and saves teams time.

3 Extend detection to calendar systems and emerging attack techniques

Calendar invite abuse emerged as one of 2025's most significant developments, with over 20% of Q4 callback phishing attacks delivered via ICS phishing. Unlike traditional email attacks, calendar invites automatically populate user calendars with malicious content depending on the organization's settings.

This gives attackers two payload delivery methods: the email itself and the calendar entry, of which the latter is not typically auto-remediated by most email security solutions. To close this gap, security teams need full visibility and control over calendar objects (create, update, delete), automated retroactive cleanup workflows, and a detection engine that treats ICS content as a significant attack surface.

4 Deploy visual and behavioral analysis beyond traditional text analysis

The rise of QR code phishing and sophisticated brand impersonation demands security capabilities that extend far beyond text analysis. Organizations should implement computer vision systems capable of detecting QR codes embedded in PDFs, images, and SVG files, then analyzing their decoded destinations through redirect chains and gating mechanisms. Natural Language Understanding must move beyond simple keyword detection to recognize manipulation tactics, urgency language, authority exploitation, and the psychological pressure patterns that define modern BEC attacks.

Thread hijacking and conversation impersonation, now comprising 28.1% of BEC attacks, require cross-message context analysis that understands email conversation patterns, recognizes when unfamiliar senders insert themselves into legitimate threads, and detects the subtle inconsistencies that distinguish authentic business discussions from fabricated ones.

5 Baseline normal communication patterns to detect anomalies

Organizations should deploy detection systems that build and maintain adaptive, privacy-aware behavioral baselines for every mailbox. Because human behavior is constantly changing, these systems must continuously evolve to tolerate organizational shifts without generating excessive noise.

With robust baselines established, detection engines can identify meaningful anomalies in real time, flagging email bombs when volume spikes or surfacing VIP impersonation attempts when messages deviate from established patterns. Look for solutions that combine automated, continuous learning with the flexibility to define custom behavioral rules for high-risk users.

 **Sublime's next-gen email security offers the AI agents, ML-powered enrichments, and advanced, real-time analysis teams need to stay ahead of attackers.**

[Get a live demo to see how →](#)

This report was compiled by the Sublime Machine Learning and Detection teams from anonymized customer telemetry and analyzed under strict adherence to our internal Data Privacy Framework.

Data

Our analysis draws from multiple complementary and anonymized data sources to provide a comprehensive view of the threat landscape:

- **Detection telemetry:** Analysis over messages processed by Sublime Security's detection engine across most enterprise customer environments from Jan. 1, 2025 to Nov. 24, 2025.
- **External research:** Correlation with findings from leading security researchers, vendors, and government agencies.

This multi-layered privacy architecture ensures that organizations can benefit from collective threat intelligence without compromising individual privacy or exposing sensitive operational details. By implementing these measures, we maintain a responsible balance between security intelligence sharing and privacy preservation.

Analysis

Our analytical approach combines automated pattern detection with expert human analysis:

- **Machine learning classification:** Initial categorization of threats using our supervised ML models trained on verified malicious and benign content.
- **Behavioral analysis:** Examination of attacker TTPs, including delivery mechanisms, social engineering techniques, and technical indicators.
- **Temporal analysis:** Tracking of threat evolution over time to identify emerging trends and anticipate future developments.
- **Cross-industry comparison:** Identification of industry-specific targeting patterns and customization techniques.

Limitations

We acknowledge several limitations in our analysis that may affect the generalizability of our findings:

- **Customer composition:** Our customer base skews toward enterprise, potentially underrepresenting threats targeted at SMBs.
- **Detection bias:** Our visibility is limited to attacks that reach our customers' email environments and does not capture threats blocked upstream by gateways or other security controls.
- **Geographic concentration:** While we have global coverage, our customer base is primarily in North America and Western Europe, potentially underrepresenting regional threats in other parts of the world.
- **Temporal window:** This report focuses on 2025, and may not capture longer-term trends.



Questions? Feedback? Reach out at research@sublimesecurity.com or [@sublime_sec](https://twitter.com/sublime_sec) on X.