STRONGESTLAYER THREAT RESEARCH

Evasion Technique Combinations in Enterprise Phishing Campaigns

December 2025 – February 2026

Author: StrongestLayer Threat Research Team

Version 1.0

February 2026

Executive Summary

This report presents findings from a comprehensive analysis of email-based threat detections collected across multiple enterprise environments between December 2025 and February 2026. Building on methodology established in our earlier Multi-Channel Evasion Research Brief (September–November 2025), this study maps how the evasion technique landscape has shifted — and hardened — in three months.

The central finding is structural. We identified **over 1,400 unique evasion technique combinations** — a 130% increase from the prior study. **34.7% of all detections now score 7+/10 on regex difficulty**, up from 20.1%. The average detection contains 4.11 evasion techniques simultaneously, up from approximately 3.2 in the prior period. This is no longer a landscape of one-trick attacks; it is a mature, layered evasion ecosystem where attackers construct multi-technique architectures specifically designed to defeat the email security stack they're targeting.

KEY FINDINGS AT A GLANCE

- **TOAD is now the single largest attack family** at 27.8% of all detections — more than 1 in 4 emails use a phone number as the payload, with no scannable technical indicator.
- **The QR Triple Stack** (QR Code + PDF Attachment + CAPTCHA Gate) has matured into a distinct attack family concentrated against Microsoft environments (QR rates up to 53%).
- **35.9% of detections fall into structural blind spots** — attack categories where detection is not merely difficult but architecturally impossible within an email-text-scanning paradigm.
- **Evasion profiles are platform-specific.** Microsoft environments face image-based payloads; Google Workspace faces notification spoofing and voicemail pivots — consistent with attackers optimizing for specific platform architectures.
- **DocuSign impersonation at 24.8%** has surpassed Microsoft as the most impersonated brand, leveraging SendGrid infrastructure that mirrors legitimate DocuSign delivery paths.

Data notes: This report intentionally excludes customer-level attribution. All technique, combination, and temporal analyses cover the full detection corpus. Where SEG-specific analyses appear, they are attributed by technology stack - not by customer - to identify differences in defensive posture across platforms. Patterns should be read as directional, not precise.

Survivorship bias: Every detection in this dataset is an email that *reached* StrongestLayer's pipeline — meaning it had already bypassed whatever upstream SEG was in place. Our technique prevalence rates represent "what gets through," not "what exists in the wild." This is a critical distinction: we are measuring the failures of email security, not the total threat landscape.

Methodology

We analyzed approximately 5,000 email-based threat detections that successfully evaded one or more Secure Email Gateways, collected from StrongestLayer's detection infrastructure over a two-month period (December 2025 – February 2026). Each detection record contained sender metadata, subject lines, and structured analysis fields covering intent, deception, anomaly, harm, and related evidence.

We enumerated 22 evasion techniques across five tactical categories — channel-shifting, URL evasion, content evasion, authentication evasion, and social engineering — and performed full-text indicator extraction with contextual validation against each detection record. Multiple extraction patterns per technique ensured comprehensive coverage; contextual requirements ensured precision. Detections were then classified into an **Attack Family Taxonomy** that groups attacks by their primary evasion strategy rather than raw technique count, capturing qualitative differences in SEG detection resistance that additive scoring alone misses.

From there, we performed clustering analysis on technique co-occurrence to identify dominant multi-technique combination patterns, mapped attributed detections to their underlying email security architectures (Microsoft and Google Workspace), and cross-referenced findings against published threat research from Proofpoint, Trustwave SpiderLabs, Sublime Security, Intel 471, and Cisco Talos to validate external consistency.

This study intentionally focuses on evasion technique analysis — understanding how these attacks made it through and identifying the dominant bypass patterns. AI authoring indicators, AI-uniqueness scoring, and temporal volume analysis have been covered in prior StrongestLayer research and are excluded from this scope. All detections in this dataset represent emails that bypassed upstream SEG detection; technique prevalence rates reflect what penetrates, not what exists in the broader threat landscape.

TOAD: The Dominant Evasion Strategy

At 27.8% of all detections, TOAD (Telephone-Oriented Attack Delivery) is the single largest attack family in our dataset. This is not a marginal channel-shift or an emerging trend — it is the primary bypass strategy for a plurality of attackers we observe. More than one in four detections include a phone number as part of the attack chain.

The mechanics are straightforward. A victim receives an email impersonating PayPal, Norton, McAfee, or Geek Squad claiming they've been charged $249.99 for a subscription renewal. The only call-to-action is a phone number. No URL. No attachment. No technically malicious content. The recipient calls the number, reaches a call center staffed by social engineers, and over the course of a 20-to-30-minute conversation is guided into credential disclosure, remote access tool installation, or gift card fraud.

Why this bypasses every email security system: the "payload" is a phone number — a string of digits indistinguishable from a legitimate business contact. A regex rule that blocks financial language combined with a phone number would fire on every legitimate billing notification in the enterprise. **This is a category of attack that operates outside the detection model email security was designed for.**

PRIOR RESEARCH

Our findings extend existing public research on TOAD. Proofpoint, who coined the term, reports 10 million TOAD attacks monthly with 67% of businesses affected. Trustwave SpiderLabs documented the same callback pattern — fake invoices with "bogus customer hotline numbers for dispute resolution." Intel 471 found that threat actors consider phone-based social engineering more reliable than email-only delivery. Cisco Talos documented PDF-based TOAD delivery as a growing vector. Sublime Security's 2026 report corroborates TOAD as a dominant trend. Our data confirms and extends these findings with evasion technique co-occurrence analysis and sub-family classification.

Our data reveals four distinct TOAD sub-families, each exploiting a different psychological vector:

Branded Financial Callback (13.2%)

The largest individual attack family in the entire dataset. Every detection combines Financial Lure + Brand Impersonation + Phone Callback. 77.9% also include Authority Impersonation, 71.1% Domain Lookalike. The attack pattern: an email impersonating a known brand (Norton, McAfee, Geek Squad, PayPal, Amazon, QuickBooks) claims a subscription renewal has been processed for a specific dollar amount, and provides a phone number to "cancel" or "dispute" the charge.

What makes this family effective is the *specificity*. The dollar amounts are oddly precise — $249.99, $349.74 — because round numbers feel fake and irregular amounts feel transactional. The order numbers ("ZJY88Q0Z", "CO5SV448") add bureaucratic credibility. Every component individually passes SEG analysis because every component appears in millions of legitimate emails.

Financial Callback Without Brand (9.1%)

Similar to the branded variant but without explicit brand impersonation — generic "payment remittance," "ACH confirmation," and "wire transfer" themes with a callback number. Counterintuitively, the absence of brand impersonation makes these *harder* for brand-matching detection rules to catch. They rely on financial urgency alone, and the false positive surface is broader: blocking "payment" + "phone number" affects far more legitimate business email than blocking "McAfee" + "renewal" + "phone number."

Notification Callback (1.1%) and Generic Callback (4.3%)

The remaining TOAD sub-families include notification-formatted callbacks — mimicking automated billing systems rather than personal messages — and generic callback variants that lack both financial and brand signals. The notification variant is particularly effective because recipients are conditioned to trust automated system notifications; the format itself conveys institutional authority.

THE TOAD TRAJECTORY

TOAD showed a ~487% increase from December to January 2026. Even accounting for a volume spike from new customer onboarding in Week 1, TOAD is growing faster than any other attack family. This is not a seasonal fluctuation — it represents a structural shift in how attackers are choosing to deliver social engineering.

The Three-Layer Evasion Architecture

The most sophisticated attacks in this dataset don't rely on any single evasion technique. They construct a layered architecture where each layer defeats a different detection capability. We observe three distinct layers:

Stage	What Happens	Why Detection Fails
Layer 1: Trusted Delivery	The email arrives through legitimate infrastructure — SendGrid, Google Calendar notifications, SharePoint sharing alerts, or redirect URLs through google.com/url endpoints.	Reputation-based filtering is defeated. The sending infrastructure and URL domains have perfect reputation scores that no SEG will block.
Layer 2: Anti-Scanner	The malicious payload is protected by CAPTCHA gates, encoded inside QR codes (images inside PDFs), or delivered as image-based payloads with no parseable text.	Sandbox analysis is defeated. The SEG follows the URL, hits a CAPTCHA, and marks it as "clean" because it only saw the challenge page, not the phishing page behind it.
Layer 3: Channel-Shift	The actual exploitation moves off the email plane entirely — to a phone call (TOAD), a mobile device (QR scan), a Teams meeting, or an SMS message.	Post-delivery remediation is defeated. Even if the email is retroactively identified as malicious, the damage occurs in a channel the SEG cannot monitor or recall.

Each layer is independently effective, but the combination creates a detection gap that compounds multiplicatively. Consider the QR Triple Stack: the QR code is an image inside a PDF (no URL text to scan), the CAPTCHA blocks the sandbox from following the decoded URL, and the QR code is scanned on a personal mobile device outside enterprise security. Three layers, three different detection capabilities defeated.

Critically, **35.9% of all detections** fall into what we classify as **structural blind spots** — attack categories where SEG detection is not merely difficult but architecturally impossible within an email-text-scanning paradigm. TOAD accounts for 27.8%, QR Code attacks for 6.0%, and Voice Channel Shift for 2.1%. An additional ~20% use CAPTCHA gates that actively detect and block SEG sandbox analysis. The remainder exploits legitimate infrastructure where blocking the delivery mechanism (Google redirects, SharePoint links) would cause unacceptable false positives.

The Evasion Technique Landscape

We track 22 individual evasion techniques across five categories. Rather than presenting them as a flat ranking, it is more useful to understand them as components of the layered architecture described above. Each technique falls into one of three roles: it either shifts the channel, obfuscates the payload, or exploits human psychology.

Channel-Shifting Techniques

These techniques move the attack payload off the email plane entirely, rendering email-only detection structurally blind. Phone Callback (TOAD) at 27.9% is now the third most prevalent individual technique in the entire dataset, behind only Authority Impersonation (65.5%) and Domain Lookalike (56.4%). QR Code at 6.0% is up from ~4.0% in the prior period. Voicemail Pivot at 2.5% was elevated to a full technique this cycle based on volume.

URL Evasion Techniques

Legitimate Redirect at 43.0% remains the most common URL-level evasion — URLs laundered through Google, Bing, LinkedIn, and Microsoft redirect endpoints that no SEG can block. CAPTCHA Gate has risen to 19.3%, approaching 1 in 5 detections. This is significant because CAPTCHA gates represent a qualitative escalation: they don't just hide the malicious signal, they *actively detect and block the detector*. Multi-hop Redirect at 19.5% chains 2+ redirect hops, often exceeding SEG recursion depth. Legitimate Service Abuse (LOTS) at 20.3% hosts malicious content on SharePoint, Google Drive, Firebase, and similar platforms with perfect reputation.

Social Engineering Techniques

Authority Impersonation at 65.5% and Domain Lookalike at 56.4% remain the bedrock of most attacks — nearly two-thirds of all detections impersonate an authority figure or organization, and more than half use a lookalike domain. Financial Lure at 44.1% is up meaningfully, consistent with the growth of TOAD's branded financial callback family. These techniques succeed because their vocabulary is identical to legitimate business communication. A regex targeting "invoice" + "payment" + "urgent" would flag every accounts payable department in the enterprise.

Technique Stacking

The average detection employs 4.11 evasion techniques. 56.8% use four or more techniques simultaneously. Only 9.1% use a single technique or none. At the extreme end, we observe detections stacking 8–11 techniques, scoring a perfect 10/10 on regex difficulty. This is a mature evasion ecosystem, not opportunistic one-technique attacks. Sublime Security's 2026 report corroborates this, finding that 34.7% of attacks employ multi-technique evasion stacking.

The Evolved Attack Family Taxonomy

Rather than scoring difficulty by raw technique count, we developed a tactical family taxonomy that classifies each detection by its *primary evasion strategy*. This captures qualitative differences that additive scoring misses. A TOAD attack with 2 techniques is harder for a SEG to detect than a redirect chain with 6, because the TOAD payload (a phone number) is structurally unblockable.

Super-Family	% of Total	SEG Detection Difficulty
TOAD / Phone Callback	27.8%	Structural blind spot
Redirect Chain Families	27.4%	Hard (evasion-dependent)
Legitimate Infrastructure (LOTS)	16.7%	Hard (reputation-laundered)
Social Engineering Pure-Play	14.0%	Medium-hard (high FP cost)
QR Code Attack Families	6.0%	Structural blind spot
Voice Channel Shift	2.1%	Structural blind spot
Content Evasion	2.0%	Medium
Collaboration Platform Spoofing	0.7%	Hard
Authentication Exploitation	0.5%	Very hard
Other/Mixed	2.7%	Varies

Table 1. Attack family taxonomy. Three of the top six families (TOAD, QR, Voice) are structural blind spots where SEG detection is architecturally impossible.

The redirect chain families (27.4%) are the second-largest group, including multi-hop chains that bounce through 2+ legitimate domains, CAPTCHA-gated phishing where automated scanners are specifically detected and blocked, and combinations of both. LOTS attacks (16.7%) host payloads on SharePoint, Google Drive, Firebase, and other platforms with perfect reputation scores — Sublime Security notes a shift toward "uncommon or emerging platforms" like Jotform, Typeform, and Notion for LOTS hosting.

Social Engineering Pure-Play (14.0%) succeeds without any technical evasion at all — purely through authority impersonation, financial urgency, and lookalike domains. These attacks contain either no URL or a URL with no malicious technical signature, relying entirely on human psychology and the difficulty of distinguishing lookalike domains at scale.

How Attackers Adapt to Your Security Stack

The most striking finding in the SEG-specific analysis is not the individual technique rates — it is the *pattern* across environments. Attackers are not deploying the same evasion strategy everywhere. They are adapting their approach based on the email security architecture of their target.

Microsoft Environments: Image-Based Payload Attacks

Across the Microsoft environments we observe, the dominant evasion strategy is image-based payload delivery. QR code rates range from 11% to 53%, with the highest concentration in environments lacking E3/E5-tier protections like Safe Links, Safe Attachments, and advanced attachment sandboxing. PDF attachment rates reach 20–43%, and CAPTCHA gate rates reach 15–38%. This is overwhelmingly the QR Triple Stack family — attackers have identified that Microsoft's native protections struggle with payloads hidden inside images inside documents.

DocuSign impersonation dominates branded attacks in Microsoft environments, reaching 60.5% in the environment with the weakest native protections. The attack pattern is specific: fake eSign requests with QR codes embedded in PDF attachments. Even in higher-tier Microsoft environments, Multi-hop Redirect penetration reaches 48% — suggesting that Safe Links URL rewriting and time-of-click protection is being systematically defeated by redirect chains that resolve to different destinations depending on when they're accessed.

Google Workspace: Notification and Voicemail Exploitation

The Google Workspace environments show a consistent, platform-specific pattern: elevated Voicemail Pivot (10–34%), high Notification Spoofing (31–59%), and high Authority Impersonation (78–82%). This is consistent with Google Workspace's deep integration of email, calendar, and notification systems. Users receive legitimate notifications from Google Calendar, Google Drive, and Google Meet constantly — making spoofed notifications psychologically plausible in a way they are not on other platforms.

Google Workspace environments also show elevated BEC rates (up to 13.9% versus a 2.6% average), CAPTCHA gate rates up to 44%, and Encoded URL rates up to 25% — suggesting attackers are layering more technically sophisticated URL obfuscation against Google's native filtering.

> **THE PLATFORM DIVIDE**
>
> Microsoft environments face image-based payload attacks — QR codes at up to 53%, PDF attachments at up to 43%. Google Workspace faces notification and voicemail exploitation — spoofed notifications at up to 59%, voicemail pivots at up to 34%. This is not random variation. The consistency of these patterns across environments suggests attackers are adapting their evasion strategies to the specific architectural surfaces of each platform, targeting the integration points that create the largest detection gaps.

SEG Penetration Rate Comparison

The following table shows evasion technique penetration rates across Microsoft and Google Workspace environments. Ranges reflect variation across multiple environments observed within each platform. Sample sizes are small; patterns are directional.

Technique	Microsoft	Google Workspace
QR Code	**11–53%**	3–14%
Phone Callback (TOAD)	9–18%	10–12%
Voicemail Pivot	7–30%	**10–34%**
Legitimate Redirect	37–60%	**37–66%**
CAPTCHA Gate	15–38%	**15–44%**
Multi-hop Redirect	**17–48%**	17–36%
Notification Spoofing	25–44%	**31–59%**
PDF Attachment	**21–43%**	5–12%
Domain Lookalike	58–66%	**66–77%**
Authority Imperson.	65–76%	**78–82%**
Financial Lure	22–23%	14–26%
Brand Impersonation	5–12%	**11–35%**

Table 2. SEG penetration rates by platform. Ranges reflect variation across environments. Bolded = platform with higher observed rates.

The DocuSign-SendGrid Nexus

DocuSign impersonation at 24.8% of all detections is the single most important brand abuse trend in this dataset, surpassing Microsoft as the top attack vector. This is not merely about volume — it is about *architectural advantage*.

Document-signing impersonation works because: DocuSign emails are expected to contain links to external documents. The "sign this document" call-to-action creates legitimate urgency. DocuSign emails are routinely sent from third-party domains via SendGrid/Twilio infrastructure, making sender domain analysis unreliable. And the DocuSign → Microsoft authentication flow ("sign the document" → "authenticate with your O365 credentials") is a workflow that enterprise users encounter regularly.

The DocuSign + SendGrid pairing deserves special attention. SendGrid is DocuSign's actual email infrastructure provider — legitimate signature requests are delivered via SendGrid. Attackers have adopted this same delivery path — using SendGrid infrastructure or SendGrid-branded sender domains to deliver DocuSign phishing, creating a combination that is indistinguishable from legitimate workflow at the email header level.

Brand Ecosystem	% of Total
DocuSign	24.8%
Microsoft	11.7%
Google	10.7%
Norton/McAfee/Geek Squad	10.6%
SendGrid	8.2%
PayPal	5.9%
Amazon/AWS	5.9%
Meta/Facebook	5.4%
Shipping (USPS/UPS/FedEx/DHL)	3.6%

Table 3. Brand impersonation prevalence. 66.2% of detections reference identifiable brands.

Campaign Intelligence: How These Attacks Work

The statistics above describe the *what*. The walkthroughs below describe the *why* — how these multi-technique combinations interact to defeat detection at each stage of the email security pipeline. Each walkthrough uses an anonymized real detection from the dataset.

The Canonical TOAD Attack

Cluster #1: Authority + Brand + Domain Lookalike + Financial Lure + Phone Callback · 3.4% of detections · 5 techniques · Difficulty 7/10

An email arrives from "Sandra" at an Outlook.com address. Subject: "welcome." The body claims to be from PayPal: "Your PayPal account has been upgraded to a business account. A charge of $249.99 has been applied. If you did not authorize this upgrade, please call (XXX) XXX-XXXX immediately to prevent identity theft."

Stage	What Happens	Why Detection Fails
Delivery	Email from webmail (Outlook.com) with PayPal branding, financial language, and a phone number.	No malicious URL or attachment to scan. PayPal branding appears in millions of legitimate emails.
Trust	PayPal logo, color scheme, legal footer. Specific amount ($249.99) and "business account upgrade" framing.	Brand-matching would require visual analysis beyond text scanning. Odd dollar amount feels transactional, not fabricated.
Exploitation	Victim calls the number. Call center performs credential harvesting, remote access installation, or gift card fraud.	The attack occurs over the phone — a channel no email security system monitors. Even retrospective email remediation cannot undo a completed phone call.

The QR Triple Stack

Family 4A: QR Code + PDF Attachment + CAPTCHA Gate · 1.5% of detections · Concentrated against Microsoft Basic (53% QR rate)

An email arrives with subject "[COMPANY] — December Bonus and Allocation for All Staff" from what appears to be an internal HR system. The body contains generic compliance language. Attached: a PDF document. Inside the PDF: a QR code. Scanning the QR code leads to a CAPTCHA-gated credential harvesting page.

Stage	What Happens	Why Detection Fails
Delivery	HR/payroll-themed email with a PDF attachment. No URL in the email body.	No malicious URL text for regex to match. PDF is a legitimate business document format. HR compliance themes are common.
Payload	QR code embedded as an image inside the PDF. Encodes a URL to a credential harvesting page.	The URL is invisible to text parsers — it exists as pixels in an image inside a document. Requires OCR + computer vision + QR decoding.
Gate	CAPTCHA (typically Cloudflare Turnstile) blocks automated analysis of the decoded URL.	The SEG sandbox, if it even decodes the QR code, hits the CAPTCHA and cannot proceed. URL is marked "clean."
Exploitation	Victim scans QR on personal mobile device, completes CAPTCHA, enters credentials on harvesting page.	Mobile device is outside enterprise security monitoring. The entire exploitation occurs on an unmanaged device.

The Multi-hop Redirect Chain

Cluster #4: Authority Impersonation + Legitimate Redirect + Multi-hop Redirect · 1.3% of detections · Highest penetration against MS E3 (48%)

An email arrives advertising weight loss medication. Embedded URLs redirect through tracksmo.store → frametwistnine.com → final credential harvesting page. Each hop individually passes reputation checks.

Why multi-hop defeats SEGs: SEG URL scanners typically follow 1–2 redirects. Multi-hop chains add 3+ hops, potentially exceeding the scanner's recursion depth. Additionally, redirect chains can be time-sensitive — the intermediate redirect returns a benign page when scanned, then redirects to the malicious page hours later when clicked. The fact that E3's Safe Links is being systematically defeated (48% penetration) suggests attackers are specifically engineering redirect chains to evade time-of-click rewriting.

The Internal Impersonation Attack

Cluster #10: Authority Impersonation + Domain Lookalike + Legitimate Redirect · 1.0% of detections

Unlike the external brand impersonation in the attacks above, this cluster impersonates the *victim's own organization*. An email arrives from dsmgroup.cl (a Chilean domain) with the subject "[COMPANY-B] Please confirm your account Verification." The sender name shows "[COMPANY-B]" — the target company's name. A "Continue" button redirects through multiple domains before arriving at navco.in, a credential harvesting page.

This is particularly insidious because it targets the trust employees have in their own company's internal systems. "Our IT department sent a verification email" is a completely plausible scenario that happens routinely in enterprise environments. The only technical signal is the sender-name vs. sender-domain mismatch, which requires contextual understanding most SEGs don't perform.

The Fully-Loaded Notification TOAD

Cluster #3: Authority + Brand + Domain Lookalike + Financial Lure + Notification Spoofing + Phone Callback · 1.8% of detections · 6 techniques

Subject: "The $349.74 USD order is currently being activated." Sender: "Segundo B" from colegiodonvasco.edu.mx — an educational institution. The email is formatted as an automated Norton billing notification with order number, transaction ID, and a phone number: "(856) 263-1091 — call if you did not authorize this charge."

The addition of Notification Spoofing is what distinguishes this from Cluster #1. When an email appears to come from an automated system, it feels like a system event rather than a human message. Recipients are conditioned to trust automated notifications because they receive hundreds from legitimate services. The fake notification format also provides cover for slightly unusual formatting — "automated systems sometimes look off." The educational domain (colegiodonvasco.edu.mx) sending a Norton billing alert is the kind of mismatch that requires sender-brand correlation intelligence most SEGs don't perform.

Where Email Security Cannot Follow

The analysis reveals three tiers of evasion that map to three distinct levels of architectural limitation in email security:

Tier 1: Structurally Impossible (35.9%)

These attacks move the malicious action off the email plane entirely. No email-centric security solution can detect, block, or remediate the actual harm because it occurs in a channel the SEG cannot monitor. TOAD (27.8%) — the phone call is the attack. QR Code (6.0%) — the URL is opened on a personal mobile device. Voice Channel Shift

(2.1%) — the victim engages with a social engineer over the phone.

Tier 2: Theoretically Detectable, Practically Defeated (~20%)

CAPTCHA-gated attacks across all families. The SEG sandbox follows the URL, encounters a CAPTCHA, and cannot proceed. The URL is marked "clean" because the scanner only saw the CAPTCHA page. This is not evasion through obscurity — it is the attacker *detecting the detector* and specifically blocking it.

Tier 3: Detectable but Prohibitive False Positive Cost (~44%)

Legitimate Redirect Abuse (43.0%) — blocking google.com redirects is not viable. Financial Lure (44.1%) — blocking "invoice" + "payment" would paralyze accounts payable. Authority Impersonation (65.5%) — the attack vocabulary is identical to legitimate business communication. These techniques succeed not because they're sophisticated, but because the cost of detecting them exceeds the cost of missing them.

Why Endpoint Detection Doesn't Close the Gap

A natural response to the structural blind spots above is: *"That's fine — our EDR will catch it downstream."* This assumption deserves scrutiny. EDR products — Microsoft Defender for Endpoint, CrowdStrike Falcon, SentinelOne — are built to detect malicious *code execution* on managed endpoints: file-based malware, process injection, shellcode, persistence mechanisms, suspicious process creation. They are effective at what they do. But the evasion patterns in this dataset are specifically engineered to avoid the endpoint entirely.

TOAD: The Attack That Never Touches the Endpoint

TOAD attacks (27.8%) deliver a phone number. The victim calls it. The social engineering — credential harvesting, remote access installation, gift card fraud — happens over the phone. EDR sees nothing: no file execution, no process creation, no network indicator of compromise. When the attacker later uses stolen credentials to access Microsoft 365, SharePoint, or Okta, the authentication happens against cloud services, not on the endpoint. From EDR's perspective, a legitimate user logged in.

The worst-case TOAD variant is guided remote access installation — the attacker talks the victim through downloading AnyDesk, TeamViewer, or ScreenConnect during the call. EDR *will* see the download and process creation. But these are signed, legitimate applications that most enterprises whitelist for IT support. Defender, Falcon, and SentinelOne will log the event but will not block it. Once the remote access tool accepts the attacker's inbound connection, the session is indistinguishable from a legitimate IT support session — the connection goes to AnyDesk's servers, not a known-bad IP.

QR Codes: The Attack on Unmanaged Devices

QR code attacks (up to 53% in Microsoft environments) are designed to move the exploitation to a personal smartphone. EDR can see the PDF opened on the corporate endpoint. It cannot see the personal phone scanning the QR code, navigating to the CAPTCHA-gated phishing page, or entering credentials. By definition, endpoint detection cannot monitor devices that don't run an endpoint agent. By the time stolen credentials are used against corporate systems, the authentication event is indistinguishable from normal use — potentially for days.

Credential Phishing: Invisible to Process Monitoring

Multi-hop redirect chains, CAPTCHA-gated phishing, and LOTS attacks all terminate in browser-based credential harvesting. EDR monitors process execution, network connections, and file system activity. It does not monitor browser DOM rendering, HTTP redirect chains (which are transparent to network-layer monitoring), or what a user types into a web form. The attacker's activity looks like normal browsing: launch browser, visit URL, enter text. No process injection, no shellcode, no malicious executable.

For LOTS attacks specifically (16.7%), the infrastructure is intentionally chosen to be trusted. EDR sees a browser connecting to s3.amazonaws.com or drive.google.com — platforms used by millions of legitimate business applications. Alerting on all S3 connections would generate tens of thousands of false positives per day. The alert would be disabled within hours.

> **THE DETECTION STACK GAP**
>
> Email security operates at the gateway layer. EDR operates at the endpoint layer. The evasion patterns in this dataset exploit the gap between them — attacks that pass through the gateway without a detectable signal, then execute through channels (phone calls, personal devices, browser interactions, cloud authentication) that the endpoint cannot monitor. As long as attackers can achieve their objectives without executing code on a managed endpoint, EDR is architecturally irrelevant to these attack patterns. The 27.8% TOAD growth documented here represents an attacker shift explicitly *away* from the endpoint.

Implications for Defense Architecture

1. URL and attachment analysis alone cannot account for the largest attack family in this dataset. TOAD at 27.8% has no URL and no attachment. The payload is a phone number, and the exploitation happens in a voice channel. Defensive architectures that assume every attack carries a scannable technical indicator will miss this category entirely — and it is the single largest category we observe.

2. Evasion strategies are adapting to specific email security architectures. The platform-specific evasion profiles we observe — image-based payloads concentrated against Microsoft, notification exploitation concentrated against Google — are consistent with attackers optimizing for the detection gaps of specific platforms. This suggests that static, platform-agnostic phishing assessments may not reflect the actual threat profile an organization faces.

3. The three-layer evasion architecture compounds across detection paradigms. Trusted delivery defeats reputation filtering. Anti-scanner technology defeats sandbox analysis. Channel-shifting defeats post-delivery remediation. Each layer is addressed by a different detection capability, and no single paradigm addresses all three. Architectures that rely on one detection approach — whether reputation, sandboxing, or content analysis — face compounding blind spots when attackers layer techniques across categories.

4. DocuSign impersonation exploits legitimate delivery infrastructure. At 24.8%, DocuSign is the most impersonated brand in this dataset. The challenge is architectural: legitimate DocuSign emails are delivered via SendGrid, so DocuSign-branded phishing sent through SendGrid infrastructure is structurally indistinguishable from real workflow at the email header level. Brand-matching rules alone cannot address this without understanding the expected delivery path.

5. For 35.9% of detections, the human recipient is the last detection surface. When attacks bypass email scanning (TOAD), avoid managed endpoints (QR codes on personal devices), and defeat sandbox analysis (CAPTCHA gates), the only remaining point where the attack can be interrupted is the human making the decision to engage. How organizations choose to support that decision point — whether through training, real-time context, or workflow controls — becomes an architectural question, not just a policy one.

Confidence Assessment

Individual technique prevalence carries high confidence — direct indicator extraction across the full corpus with cross-validation patterns. The Attack Family Taxonomy carries medium-high confidence; the classification hierarchy is applied consistently but boundary cases between families exist. SEG-specific analysis is directional; platform-level patterns are consistent across environments but sample sizes per architecture limit statistical precision. TOAD categorization carries high confidence — phone callback identification is unambiguous and externally validated against Proofpoint, Trustwave, and Intel 471 research.

Sources

- Proofpoint — "The 411 on Call Center Scams: TOAD Attack Delivery" and "Typical Attack Sequence of TOAD Threats"
- Trustwave SpiderLabs — "Hooked by the Call: A Deep Dive into Callback Phishing Tricks" (October 2024)
- Intel 471 — "To Deliver Malware, Attackers Use the Phone"
- Sublime Security — "Key Findings from the 2026 Email Threat Research Report"
- Keepnet Labs — "TOAD Explained & Defense"
- Cisco Talos / Omid Mirzaei — PDF-based TOAD delivery research (2025)
- CyberPress — "Cybercriminals Use Microsoft Entra Invitations to Deliver TOAD Attacks"
- The Hacker News — "Hackers Using PDFs to Impersonate Microsoft, DocuSign in Callback Phishing"
- Deepstrike — "Vishing Statistics 2025: AI Deepfakes & $40B Voice Scam Surge"
- Kymatio — "Phishing Trends 2026: AI-Phishing, QRishing & Voice Deepfakes"

———