

2026 Attack Landscape Report

Nearly 800,000 Attacks
Reveal How Threat Actors
Tailor Tactics to Their Targets

Thank you for your interest
in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider**® supporting a broad portfolio of industry-leading technologies through GSA Multiple Award Schedule (MAS), NASPO ValuePoint, The Quilt and a wide range of other contract vehicles.

As the **Master Government Aggregator**®, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with Abnormal AI, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.



To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com



Explore More Resources:
carah.io/abnormalresources



Join Events & Webinars:
carah.io/abnormalevents



Discover Technology Solutions:
carah.io/abnormal



Learn About Procurement:
carah.io/abnormalcontracts



Connect With Our Team:
AbnormalAI@carahsoft.com
844-214-4790

2026 Attack Landscape Report:

**Nearly 800,000 Attacks Reveal How
Threat Actors Tailor Tactics to Their Targets**

Executive Summary

Email attacks are often discussed as if they arrive at random—a numbers game played at scale, where volume alone determines who gets hit. The data in this report tells a different story. Across 159 million attacks observed during the second half of 2025¹, the clearest pattern isn't about how many attacks organizations face. It's about how precisely those attacks are tailored to their targets.



Phishing Techniques Are Calibrated to the Environment

Phishing accounts for 58% of all attacks—the largest category by volume and the one most employees will encounter. What makes phishing worth examining closely is how threat actors adapt their techniques to context.

Obfuscation and evasion methods adjust to the defensive infrastructure threat actors expect to encounter, varying by organization size. File-sharing phishing lures concentrate in industries and roles where document exchange is constant and expected. Brand impersonation scales with the complexity of an organization's software footprint. In each case, the lure is calibrated to blend into the workflows and tools employees already use. The operational characteristics that make an organization productive are often the same ones that make it vulnerable.



BEC Changes Shape as Organizations Grow

While business email compromise (BEC) represents roughly 11% of attacks by volume, the potential impact per success is far greater. Further, workforce size changes not just the volume of BEC an organization faces, but also what kind.

The data reveals a substitution effect between VIP and employee impersonation that reflects how authority and communication norms change with scale. VIP impersonation accounts for 43% of named identity impersonation at small organizations but just 7% at large enterprises. Lateral attacks—sent from compromised internal accounts—follow their own scaling logic, rising from less than 1% of BEC at small organizations to more than 23% at large enterprises, emerging only where the identity surface justifies the investment.



VEC Tactics Vary by Geography

Vendor email compromise makes up roughly 61% of all BEC and introduces a geographic dimension to targeting.

Invoice fraud dominates in North America at 42% of campaigns, while procurement-stage pretexts lead in EMEA at 41%—a divergence that maps to regional business practices. Across all four pretexts, the choice between vendor impersonation and compromise follows cost-benefit logic. Billing account updates carry the highest compromise rate at 26.5%, reflecting the scrutiny the request demands. Invoice fraud sits at the opposite extreme at under 1%, where impersonation alone gets the job done.



This report examines how modern email threats adapt to institutional context—and how they exploit it. The same structures, workflows, and relationships that define how an organization operates also define where an attack can blend in undetected.²

Threat actors modify their tactics to fit the target's operational profile, but the end goal is consistent: make malicious activity indistinguishable from routine business. The data that follows traces both dynamics: why attack methods change by target, and how they're engineered to disappear into the environment they enter.

¹ The analysis in this report is based on a statistically representative sample of approximately 797,000 messages drawn from this population. Full sampling methodology is detailed in the appendix.

² Demographic dimensions such as industry, organization size, and job function are used as analytical lenses to examine how attack characteristics vary within segments—not to rank which segments are most targeted. The appendix details the sampling approach, classification hierarchy, and composition bias guardrails that govern every finding.




Table of Contents

Phishing	04
▪ Redirects and Link Shorteners	05
▪ File-Sharing Phishing	06
▪ Brand Impersonation	07
▪ Real-World Example of Phishing	08
Business Email Compromise	10
▪ Internal Impersonation BEC	11
▪ Named Identity Impersonation	13
▪ Generic Internal Impersonation	15
▪ Lateral BEC	16
▪ Industry Insight: Higher Education's Lateral Attack Challenge	17
▪ Real-World Example of Business Email Compromise	18
Vendor Email Compromise	19
▪ VEC as a Share of BEC	20
▪ Invoice Inquiry	22
▪ Billing Account Update	23
▪ Request for Quote (RFQ)	24
▪ Payment Inquiry	24
▪ Regional Threat Profiles	25
▪ Real-World Example of Vendor Email Compromise	27
Conclusion	29
Appendix	30
About Abnormal	32



Phishing:

Credential Theft via Obfuscation, Appropriation,
and Impersonation



22%

Phishing utilizing redirect chains to
obscure malicious destinations

Phishing is the most common threat employees will encounter, accounting for 58% of all observed attacks. Defending against that volume is a challenge on its own, but the issue is exacerbated by how precisely the tactics adapt to the target environment.

12%

Credential theft attempts disguised
as document-sharing notifications


The sections that follow examine three dimensions of that adaptation. The utilization of redirect chains and link shorteners reveals how attackers calibrate their evasion techniques to the security infrastructure they expect to face, with usage patterns that shift meaningfully across organization size. File-sharing phishing exploits the routine mechanics of cloud-based document exchange, primarily targeting industries and job functions where clicking a shared document notification is an unremarkable part of the workday. Brand impersonation borrows credibility from the software platforms employees already trust, and its prevalence tracks closely with the size and complexity of an organization's technology footprint.

12%

Rate of brand impersonation
in phishing attacks

Across all three, a consistent pattern emerges: phishing techniques cluster where they're most likely to blend in. Attackers match their lures to the workflows, tools, and communication norms of the environment they're targeting—opting to exploit routine rather than attempt to circumvent it.





Top Link Shortener Domains

Domain	% of Link Shortener Usage
tinyurl.com	31.6%
t.co	26.6%
shorturl.at	10.9%
is.gd	9.0%
bit.ly	6.9%

The prominence of `t[.]co` stands out for a different reason. Unlike the other domains on this list, `t[.]co` isn't an independent shortener service; it's Twitter/X's redirect infrastructure, automatically applied to any link posted on the platform. Threat actors likely post malicious links on Twitter/X specifically to generate `t[.]co`-shortened URLs, exploiting the fact that security tools are reluctant to block a widely trusted domain wholesale. The attacker gets a clean, reputable-looking URL without registering anything.

Smaller organizations often lack sophisticated URL inspection, leaving basic redirect chains effective on their own—no additional obfuscation needed. Larger enterprises are more likely to deploy link-reputation and URL-scanning tools that would catch standard redirects, which renders shorteners especially valuable as an additional layer of obfuscation. Threat actors aren't using link shorteners indiscriminately; they appear to use them where the defensive environment demands it.



File-Sharing Phishing



File-sharing phishing is an attack in which a threat actor poses as a colleague or familiar file-hosting or e-signature service and sends a malicious link disguised as a shared document. Using services like SharePoint, Dropbox, Google Drive, or DocuSign as cover, these attacks either impersonate a legitimate platform or exploit the platform itself to deliver the email or link.

The lure is inherently low-suspicion in any environment where cloud-based document exchange is a standard workflow. But the rate isn't uniform. File-sharing phishing accounts for 12.4% of all phishing and concentrates heavily in industries and job functions where external document exchange is constant and expected.

Industries Reliant on Document Exchange

Within financial services, 22.2% of phishing attacks use file-sharing lures—nearly double the 12.4% sample average. The construction and engineering industry runs close behind at 21.3%.

The financial services industry runs on documents: loan agreements, account statements, audit packages, compliance disclosures, investment reports, etc. Receiving a notification that someone has shared a document is entirely unremarkable for a financial services employee, and attackers exploit that normalcy. A fake DocuSign request or spoofed SharePoint notification lands in a context where such communications arrive constantly and are expected to require a click.

Construction projects also generate a relentless volume of shared documents across a wide web of parties—general contractors, subcontractors, architects, engineers, project owners, inspectors—who exchange drawings, specifications, RFIs, submittals, change orders, and bid packages throughout the project lifecycle. Cloud file-sharing platforms are standard infrastructure for this workflow, meaning a “new document shared with you” notification is completely routine, especially from an unfamiliar party.

The Roles on the Receiving End

The industry view tells part of the story, but the job-category data sharpens it. The rate of file-sharing phishing skyrockets for roles whose daily workflows revolve around external document exchange. Finance and accounting leads at 25.1%—roughly double the sample average of 12.4%—followed by legal and compliance at 21.4% and sales and business development at 17.4%.

These three types of roles share a common trait: a high volume of inbound documents from external parties, delivered via the very platforms that file-sharing phishing impersonates. Finance professionals process invoices, purchase orders, and financial statements that routinely arrive via SharePoint, Google Drive, or Docusign. Legal teams receive contracts, discovery materials, and regulatory filings the same way. Sales operates on a slightly different dynamic. The documents are similar (e.g., proposals, RFPs, partnership agreements), but the senders are often unfamiliar. The entire function is built around engaging with new external parties, which means a shared-document notification from an unknown name looks like business-as-usual rather than a warning sign.

In each case, clicking a document-sharing notification is simply part of the job. The lure works because it's indistinguishable from legitimate workflows, and threat actors are choosing their pretexts accordingly.



Brand Impersonation



Across the full sample, 12% of phishing attacks involve brand impersonation—leveraging the name and visual identity of a trusted company to make a credential harvesting attempt appear as a routine notification. The tactic works by borrowing trust the recipient already extends to a familiar brand. But the rate varies significantly depending on how many branded platforms an organization's employees interact with daily and how deeply those tools are embedded in standard operations.

The Software Stack Effect

Brand impersonation is involved in 16.3% of phishing attacks targeting large enterprises, the highest of any organization size and well above the sample average. The most likely explanation is that brand impersonation scales with the enterprise software footprint.

Large enterprises are, by definition, large software buyers. They run Microsoft 365 at scale, authenticate through Okta or Microsoft Entra ID, sign documents via DocuSign, manage procurement through SAP or Coupa, collaborate in Salesforce, and transfer files through SharePoint or Box. Impersonating any of these platforms is a credible pretext precisely because employees at large enterprises genuinely use all of them, often daily. A fake MFA prompt from Microsoft or a spoofed DocuSign signature request doesn't read as suspicious—it reads as Tuesday.

Smaller organizations tend to run leaner, less standardized software stacks, which narrows the attacker's impersonation surface. A fake Okta prompt loses its power if the target organization doesn't use Okta, and the threat actor often has no way to know.

Why Hospitality Leads in Brand Impersonation

In the hospitality industry, nearly one in four phishing attacks (24.1%) feature brand impersonation—more than double the sample average. The next-closest industry is technology at 16.1%, followed by education (14.4%), advertising and marketing (13.1%), and financial services (13.1%). Healthcare sits at just 7.1%—about a third of the rate in hospitality.

The hospitality industry's heavy reliance on branded third-party platforms (reservation systems, payment gateways, review sites) creates a target-rich environment for brand impersonation. A convincing fake notification from Booking.com, Square, or a hotel loyalty program is a natural fit for an environment where such communications are routine and expected. Other industries also depend on well-known platforms, but few match hospitality's variety. A single hotel property might interact with a dozen branded services daily across booking, payment, staffing, and guest communication—all of which lend themselves to impersonation.



Real-World Example of Phishing

From: "Amex" <communications@comcast.com>
Sent: 12/31/2025 7:58:28 AM
To: "Recipients" <communications@comcast.com>
Cc:
Subject: View and Complete: Disputed Payment Posted To Your Account.

Your due date is approaching.



Dear Card Member,

A Disputed Payment Received
 Disputed Payment Posted To Your Account

We have adjusted your payment options to reflect a disputed amount of \$1218.16. to your card account
 Follow the prompt below to view chargeback status

[Complete Charge Back Payment](#)

Notice: Payment will be posted into your account within 24 hours after validation.
 Thank you for your Card Membership,

American Express Customer Care

Explore more with your American Express® Card



[Log in Online](#)



[Get the Amex App](#)



[Enroll in AutoPay](#)

DON'T do business WITHOUT IT™

[Privacy statements](#)

[Contact us](#)

[Update your email address](#)

Your account information is included above to help you recognize this as a customer care e-mail from American Express. To learn more about e-mail security or report a suspicious e-mail, please visit us at americanexpress.com/phishing. We kindly ask you not to reply to this e-mail but instead contact us via [Customer Care](#).

© 2025 American Express. All rights reserved.

SPB0FYI089

This campaign showcases how modern phishers blend flawless visual spoofing with subtle infrastructure tricks to evade traditional defenses. The email is a near-pixel-perfect clone of a genuine American Express customer care notification. Threat actors lift real Amex assets and branding directly from Amex-controlled domains and drop them into a complex, multi-section HTML template complete with a nav bar, hero banner, primary CTA, and legal footer. To both users and many filters, it looks indistinguishable from an authentic email.

The copy is carefully crafted for maximum pressure with minimal suspicion. Recipients are told that a disputed payment of \$1,218.16 has been posted, that their payment options have been adjusted, and that funds will move within 24 hours after validation. The copy frames the next step—follow the prompt to “view chargeback status” by clicking “Complete Charge Back Payment”—as a protective action, rather than something more overtly suspicious. The combination of a concrete dollar amount, tight timeline, and clear defensive action—wrapped in pristine branding—pushes users to click before they think.



The technical implementation is just as strategic. The message is sent from `communications@comcast[.]com`—a real consumer ISP domain rather than American Express itself. Because many security stacks enforce SPF and DMARC strictly for their own domains, authentication failures on this third-party sender domain are often treated as weak signals instead of hard stops, giving spoofed traffic room to reach users.

The main CTA doesn't go straight to a shady domain; it first resolves to a URL hosted on `t[.]co`, Twitter/X's official URL shortener, a high-reputation service widely allow-listed by SEGs and web proxies. From there, targets are redirected to a phishing page that reproduces the American Express online experience. At that point, any credentials entered are delivered straight to the attacker.

AMERICAN EXPRESS My Account Cards Banking Travel Rewards & Benefits Business Customer Service **Log In**

FDIC FDIC-Insured - Backed by the full faith and credit of the U.S. Government

Log In to My Account

User ID

Password

Cards and Banking

☐ Remember Me

Login

[Forgot User ID or Password?](#)

[Create New Online Account](#)

[Confirm Card Received](#)

[Visit Our Security Center](#)

Access your account on the go.

Download the **Amex® App** to manage your Card account wherever, whenever.


Terms Apply

Download Now

Because the rest of the links and assets resolve to legitimate Amex infrastructure, security engines see a message dominated by known-good domains plus one obfuscated URL. That mix heavily biases URL-reputation and static sandboxing toward a benign verdict—right up until the target hands over their username and password.

Business Email Compromise:

Deception Dictated by the Org Chart



39%

BEC that exploits trust in colleagues, executives, and internal departments

45%

Internal impersonation BEC using non-executive employee identities

13%

BEC sent from compromised internal accounts

Business email compromise (BEC) represents roughly 11% of attacks by volume—a fraction of phishing’s 58%. But the comparison understates the threat.

While BEC attacks are lower in volume, they are higher in investment per attempt, with each one built around a specific identity, a tailored pretext, and enough contextual detail to prompt action without triggering verification. The damage per success is also far greater, with the average BEC incident costing a business \$123,005, according to the FBI IC3.

This chapter focuses on the 39% of BEC that exploits trust within the target organization itself—what we call internal impersonation BEC. The remaining 61%, involving vendor and partner impersonation, is covered in the dedicated VEC chapter that follows.


Within internal impersonation BEC, four distinct tactics compete for share: employee impersonation, VIP impersonation, generic internal impersonation, and lateral attacks from compromised accounts. Organization size, industry, and recipient role each reshape the mix, since each variable changes which identities are available to exploit and which requests are plausible enough to succeed. A 500-person company and a 50,000-person enterprise face fundamentally different versions of the BEC problem—not just in volume, but in kind.



Internal Impersonation BEC



Internal Impersonation BEC by Type



Nearly 40% of all BEC attacks exploit the trust employees place in colleagues, executives, and internal departments.

Within this category, employee impersonation is the most common tactic at 45.3%, covering attacks that impersonate a named, non-executive colleague. Generic internal impersonation follows at 36.7%. These are department-level lures rather than individual ones—e.g., the fake IT helpdesk notice, the HR benefits update, the payroll system alert. VIP and executive impersonation accounts for 8.4%, despite its outsized reputation as the defining BEC tactic. Lateral attacks³—originating from genuinely compromised internal accounts rather than impersonating one—account for 9.6%.


Each represents a distinct attacker approach to the same underlying problem: how to make a request credible enough that the recipient acts on it without verifying. The sections that follow examine how those approaches vary across organization size, industry, and recipient role.

³ “Lateral attacks” refers to all lateral attacks in which the identity of the impersonated party was not internally categorized.



How Organization Size Reshapes the BEC Mix

Internal Impersonation BEC by Organization Size



The internal impersonation BEC mix at a 500-person company looks nothing like the mix at a 50,000-person enterprise. As organization size increases, VIP impersonation gives way to employee impersonation in a near-perfect inverse relationship. Alongside that shift, lateral attacks—nearly absent at small organizations—emerge as a significant share of BEC at the enterprise end.

Threat actors aren't applying a single playbook uniformly. They're adapting to the institutional realities of their target: how authority flows, how many identities are available to exploit, whether a department-level lure or a named individual is more credible, and whether compromising an internal account is worth the investment. Workforce size doesn't just scale attack volume; it fundamentally changes the types of attacks a business faces.

In large enterprises, the opposite is true. CEOs don't email the finance department directly. Multi-person approval workflows are standard, out-of-band verification is expected, and security training has made C-suite impersonation a well-known red flag. But the underlying attack still works; it just requires a different cast. A message from a peer or mid-level colleague is far less likely to trigger skepticism than one purportedly from the C-suite.

Both tactics exploit the same mechanism: the credibility of a known, named individual within the organization. The variable is simply which individual is most convincing to impersonate, given how the target organization operates.

VIP Impersonation in Nonprofits

The org-size data shows how structural factors shape which identities attackers choose to impersonate. That logic plays out at the industry level too, and nonprofits are the clearest example. Among internal impersonation BEC attacks, 42% of those targeting nonprofits involve VIP impersonation—roughly 5x the sample average of 8.4% and nearly double the next-highest industry, technology, at 24.1%. The signal is directionally strong, though the sample base for nonprofits is relatively thin.

Nonprofit leadership is often composed of high-profile public figures: executive directors, board chairs, celebrities, and major donors. Their identities are well-documented in annual reports, donor lists, and public bios, giving threat actors ready-made impersonation material. Those same public sources also reveal the organization's hierarchy, allowing attackers to easily identify who reports to whom and where financial authority sits.

The operational environment of most nonprofits amplifies the exposure. Nonprofits tend to run on lean budgets, which translates to limited access to security tooling and training. Hierarchies are often flat, giving staff more direct access to finance functions. And a mission-driven culture means urgency from leadership is expected rather than questioned—when the executive director asks for something fast, people move.

Higher turnover and reliance on part-time or volunteer staff further deepen the vulnerability. Employees less embedded in the organization are less likely to recognize when a request breaks the pattern.

Executives as Subject and Target

More than 41% of internal impersonation BEC reaching executive leadership involves VIP impersonation—by far the highest of any job category, and 5x the sample average of 8.4%. The dynamic is worth naming explicitly: executives are both the most common subjects of VIP impersonation (attackers impersonate them to reach other departments) and the most common recipients of it (attackers impersonate other executives to reach them).

But VIP impersonation works differently depending on where it lands. When a CEO impersonation arrives in a finance recipient's inbox, it works because of authority—the executive's position is the lever. When it arrives in another executive's inbox, the mechanism is different. Executives routinely work with other executives; a message from a peer or superior is simply normal communication. At this level, VIP impersonation functions less as an authority exploit and more as a peer familiarity exploit—structurally closer to employee impersonation than the name suggests.

The impersonation is still of a named, high-status individual, but the credibility comes from the relationship, not the rank. This peer-impersonation pattern within the executive layer is a distinct attack surface.



Generic Internal Impersonation



Named identity impersonation needs a convincing person; generic internal impersonation just needs a convincing department. These are the fake IT helpdesk notices and HR benefits updates that borrow authority from a function rather than an individual. The tactic accounts for 37% of internal impersonation BEC and succeeds because employees are conditioned to act on communications from internal systems without scrutinizing who actually sent them.

Lures That Match the Workflow

Of internal impersonation BEC reaching IT and technology recipients, 66.6% involves generic internal impersonation—well above the 36.7% sample average and the second-highest among named job categories after finance and accounting (72.8%).

The elevated rate is logical when you consider what generic internal impersonation actually looks like in practice: fake IT helpdesk notices, system alerts, credential reset requests, MFA re-enrollment prompts, and access provisioning emails. These are communications that IT staff receive legitimately and routinely in the course of their role. A threat actor impersonating “IT Security” and asking a recipient to verify their credentials is a far more contextually appropriate pretext when directed at someone in IT than at someone in Sales—the lure fits the workflow. The job function that is most familiar with fake helpdesk attacks is also the one for which fake helpdesk attacks are the most believable pretexts.

Finance and accounting’s even higher rate reflects the same logic applied to a different set of lures. The dominant flavor of generic internal impersonation is almost certainly different there—fake HR payroll notices, finance system alerts, procurement approvals—rather than IT helpdesk impersonation. The category covers a range of department-level lures, and the variety that dominates in a given segment tends to mirror the recipient’s actual workflow.



Lateral BEC



Lateral BEC accounts for 13% of all business email compromise, but unlike every other tactic in this chapter, the attacker isn't pretending to be someone inside the organization. They're operating from a compromised account belonging to a real employee. The message comes from a legitimate internal address, passes authentication checks, and carries the implicit trust that employees extend to emails from known coworkers. That makes lateral BEC harder to detect by design. And as organization size increases, both the opportunity and the incentive to pursue it grow dramatically.

Why Lateral BEC Scales With Size

Because "lateral" is an attack attribute rather than a standalone impersonation category, it can co-occur with other impersonation types. To capture the full scope, the rates in this section reflect lateral BEC as a share of all BEC rather than just one subcategory.

Measured this way, lateral BEC scales dramatically with organization size. At small organizations, it accounts for just 0.24% of BEC—nearly nonexistent. The rate climbs to 1.1% at mid-market organizations and 7.1% at large mid-market. Within enterprise organizations, lateral BEC jumps to 17.8%, and at large enterprises, 23.2%—nearly a quarter of all BEC, against a sample average of 13%.

Larger organizations maintain tens of thousands of email accounts, and with that scale comes a more complex identity surface—shared credentials, third-party integrations, interconnected systems—all of which create more entry points for an attacker to gain access to a legitimate account in the first place.

Once inside, the returns compound. A single compromised account in a 50,000-person organization has access to a vast pool of trusted recipients spanning interconnected departments, and internal emails move freely in ways that external messages cannot. High email volume provides additional cover: thousands of internal messages flow through inboxes daily, making a lateral attack far harder to spot than it would be in a 100-person company, where any unexpected internal message would stand out.

For threat actors, this is a straightforward ROI calculation. Compromising an internal account requires meaningful upfront investment—reconnaissance, credential theft, maintaining persistent access—but the payoff in a large organization justifies the effort. In a small organization, that same investment yields far less. Employees tend to know their colleagues personally, verification is as simple as walking over to someone's desk, and the target surface is a fraction of the size.



Industry Insight: Higher Education's Lateral Attack Challenge

Education is the dominant outlier in lateral attack data, and the pattern holds across both phishing and BEC.

Lateral phishing is rare in most industries, with just 2.3% of phishing overall originating from compromised internal accounts. In education, that rate is 7.1%—more than 7x the next-closest industry. For students, it climbs higher still: nearly one in eight phishing attacks they receive comes from inside their own institution.

The lateral BEC picture is even more pronounced, as 33% of all BEC targeting higher education is lateral. When BEC targets students specifically, that figure jumps to 54%. Across all attack types, students show a 14.9% lateral attack rate, the highest of any job category by a wide margin, against a sample average of 2.7%.

Education's lateral problem is structurally different from what lateral attacks look like elsewhere. In most industries, lateral BEC involves a compromised employee account being used to conduct financial fraud against other parts of the organization. In higher education, the pattern is peer-to-peer: compromised student accounts are used to attack other students. The cycle is self-propagating; once one account is taken over, it becomes a trusted launchpad against the rest of the student population.

Higher Education's Unique Vulnerabilities

Imagine if every four months, your security infrastructure had to accommodate thousands of new users. That's the reality of a university campus, and several features of this environment make it uniquely hospitable to lateral spread.

The entry points are abundant. Student populations are large, and password reuse is common, creating easy targets for credential harvesting. Dormant accounts—belonging to alumni, transfers, or withdrawn students—often remain active long after the owner has stopped checking them, providing unattended footholds for compromise. And unlike a corporate environment where IT can enforce password policies and decommission inactive accounts on a predictable schedule, universities manage a population that is constantly turning over, with account hygiene that varies widely from student to student.

Once inside, the attacker benefits from a trust culture built around .edu addresses, with emails from these accounts carrying implicit credibility. Academic email systems are open and federated by design, and the culture tends to prioritize accessibility over restrictive security controls. That limits the friction that would otherwise slow lateral movement. Student accounts also typically receive less monitoring and security investment than staff accounts, giving threat actors more room to operate unnoticed.


Because users constantly cycle in and out of the system, these conditions don't improve over time—they reset. Every semester brings a new cohort with new credentials, new devices, and no institutional memory of last semester's attacks.

Two Lateral Vectors, Two Different Threats

It's worth distinguishing what lateral attacks actually accomplish in this environment. Lateral phishing uses compromised accounts to deliver credential-theft attempts, expanding the pool of compromised accounts and sustaining the cycle. Lateral BEC uses them to conduct financial fraud, such as gift card and fake job scams. One is about spreading access; the other is about monetizing it.



Real-World Example of Business Email Compromise



The attack begins with an email that looks utterly routine. The Managing Director/CEO writes in German⁴ to an employee with the subject line “Required: Comprehensive Vendor Invoice Status Report” and asks for help compiling a comprehensive list of outstanding vendor invoices that are currently due.

The message is entirely text-based, with no links or attachments, and uses a measured but insistent tone. Rather than manufacturing a crisis or attempting to redirect a payment, the attacker requests structured data: vendor names and email addresses, invoice numbers, dates, total amounts, due dates, and planned payment dates.

It’s a simple, polite request framed as supporting cash-flow approval and liquidity planning, with a “Sent from my iPhone” footer to normalize brevity and any formatting oddities. That business context, plus the request to keep an internal Excel tracking sheet up to date, grounds the email firmly in familiar AP workflows.

On the technical side, the email originates from an attacker-controlled domain, hosted on their own mail server, and configured to pass authentication. SPF and DMARC both pass, and there are no signs of spoofing the recipient’s domain. This is a straight, authenticated send from external infrastructure into the target’s Microsoft 365 environment. As a result, the native Microsoft stack classifies the message as clean, non-bulk, and non-suspicious—allowing a high-quality BEC reconnaissance email to land directly in the inbox.


If the recipient complies, the attacker walks away with everything needed to execute targeted follow-on fraud: who the company pays, how much, when payments are scheduled, and which invoices are most time-sensitive. With that information, an adversary can craft highly targeted follow-on BEC messages—impersonating vendors or executives to redirect upcoming payments—with a level of specificity that traditional email security is poorly equipped to detect or stop.

⁴ The email has been translated into English for illustrative purposes.



Vendor Email Compromise:

Financial Fraud Engineered for the Target Market



61%

BEC involving vendor or partner impersonation

The prior chapter examined business email compromise (BEC) attacks that exploit trust between colleagues within the same organization. This chapter turns to the larger share of the problem: vendor email compromise (VEC), which accounts for the majority of all BEC.

41%

High-risk VEC campaigns using invoice inquiry pretext

What makes VEC especially difficult to defend against is that billing and payments are a routine part of the vendor-customer relationship, discussed over email every day. Consequently, malicious messages seemingly from vendors requesting changes to banking information or large fund transfers may not be immediately flagged as suspicious.

The scale of modern supply chains compounds the problem. Most organizations work with dozens or hundreds of vendors, and no single employee can be fluent in every vendor's communication patterns, invoicing norms, or usual points of contact. Threat actors exploit that gap, using spoofed sender addresses, lookalike domains, and in some cases genuinely compromised vendor accounts to make their messages indistinguishable from legitimate correspondence.

81%

Vendor-related BEC targeting shared mailboxes


Four pretexts account for the full distribution, but they differ sharply in the level of credibility they demand—and attackers adjust their techniques accordingly. Those patterns also shift by region, reflecting differences in procurement norms that threat actors appear to understand and exploit.



VEC as a Share of BEC



Breakdown of BEC by Attack Category




Approximately 61% of all BEC involves the impersonation of a vendor or business partner rather than an internal employee, executive, or department. Within that share, two tiers emerge: 26% classified as high-risk VEC and 34.9% as standard VEC.

The distinction between the two tiers reflects analytical depth, not threat severity. High-risk VEC campaigns are those with verified categorical labels that enable detailed analysis of attack pretexts, impersonation techniques, and targeting patterns. Standard VEC captures the broader universe of vendor and partner impersonation that falls outside that verified set.

The sections that follow draw on a dedicated VEC dataset—a census of all high-risk VEC campaigns during the period, totaling 18,500+ campaigns. This separate sample powers the granular analysis of attack pretexts, technique selection, and regional variation.



VEC Pretexts and Techniques



Among 18,500+ high-risk VEC campaigns, four attack pretexts account for the full distribution: invoice inquiry (41.3%), billing account update (23.6%), payment inquiry (19.8%), and RFQ (15.3%). Across all four, the vast majority of attacks use impersonation—an attacker creating a lookalike email address or spoofing a known vendor—rather than compromising the vendor’s actual account. Overall, 87.5% of high-risk VEC is impersonation-based, while just 8.95% involves a genuinely compromised vendor account. But that top-line figure obscures a more revealing pattern. The compromise rate isn’t uniform across pretexts; it tracks closely with the amount of credibility the pretext requires to succeed.

Invoice inquiry sits at one extreme: 98.1% impersonation, 0.95% compromise. Sending a plausible-looking invoice from a lookalike domain requires almost no access to the real vendor’s systems. The lure is familiar, the volume of legitimate invoices creates natural cover, and the recipient has no strong prior reason to verify the sender’s authenticity. When impersonation gets the job done, there’s no need to go further.

Billing account updates sit at the other extreme: 73.2% impersonation, 26.5% compromise—by far the highest compromise rate of any pretext. The request itself explains the gap. Asking a company to change the bank account to which it routes vendor payments is inherently suspicious; most organizations with any financial controls will scrutinize it. A lookalike email address may not survive that scrutiny. A message from the vendor’s actual account—carrying the implicit legitimacy of a real business relationship—is a different proposition entirely. When the pretext demands credibility, threat actors invest in acquiring it.

RFQ (78.8% impersonation, 12.8% compromise) and payment inquiry (89% impersonation, 1.7% compromise) fill in the middle of the spectrum in predictable ways. RFQs require enough credibility to initiate a new vendor relationship, but not enough to redirect existing payment routing. Payment inquiries are broad, low-investment probes—more about finding a responsive target than executing a specific fraud—and their low compromise rate reflects that minimal effort.

The throughline is straightforward: attacker technique selection follows cost-benefit logic. Impersonation is the default, and compromise is the upgrade, deployed where impersonation alone won’t close the gap.

Invoice Inquiry



Invoice inquiry—the submission of a fraudulent or manipulated invoice designed to trigger a payment to an attacker-controlled account—is the most common VEC pretext. The approach is high-volume and low-effort by design, as a convincing fake invoice from a lookalike domain requires no prior relationship and no access to the real vendor’s systems. That simplicity makes it broadly effective, but it lands hardest in industries, roles, and mailbox types where processing invoices is constant, high-volume work.

Invoice Fraud in Media and Entertainment

Of high-risk VEC campaigns targeting media and entertainment, 64.7% use the invoice inquiry pretext—the highest among for-profit industries—while the overall account compromise rate is just 3.7%, well below the sample average of 9%. Both figures point in the same direction.

The pattern maps to how the industry operates. Production companies, studios, publishers, and media agencies run on a large, fluid ecosystem of freelancers, contractors, and production vendors whose relationships begin, end, and rotate constantly. An invoice sent from an unfamiliar name or a newly registered domain isn't inherently suspicious in an environment where working with new vendors is the norm. Threat actors don't need to compromise real vendor accounts to execute invoice fraud against media companies. They send plausible-looking invoices from lookalike domains because the environment doesn't require more.

Where Invoice Fraud Lands

Finance and accounting recipients see invoice inquiry as the dominant VEC pretext at 56.7%—approximately 1.4x the sample average of 41.3%. Executive leadership is close behind at 51.5%. The concentration follows the workflow: processing invoices is a core function of these roles. Finance teams receive them at high volume from a wide range of vendors, often under payment-term deadlines that create pressure to act quickly. A fraudulent invoice landing in that queue doesn't need to be particularly sophisticated; it just needs to look like one more item in an already full pipeline.

The near-equivalent rate at executive leadership points to a complementary dynamic. Executives don't process invoices themselves; they approve them. And attackers likely understand that an invoice forwarded with an executive's implicit endorsement can accelerate payment and short-circuit the verification steps that might otherwise catch it.


The Shared Mailbox Vulnerability

More than four out of five BEC attacks targeting shared mailboxes are vendor fraud. When high-risk and standard VEC are combined, 81% of BEC to shared mailboxes is vendor-related—more than any other recipient category. Within those VEC attacks, invoice fraud dominates: 56.6% of VEC campaigns targeting shared mailboxes use an invoice inquiry pretext.

Shared mailboxes like accounts-payable@, invoices@, and purchasing@ are perfectly positioned for invoice fraud. They exist purely to process payments and handle financial correspondence at volume, creating ample cover for fake invoices to hide among legitimate ones. Additionally, the addresses are often publicly discoverable or easily inferred from naming conventions. Because multiple users have access to them, employees also have fewer opportunities to understand the relationship context that would help them spot anomalies. There's no "this doesn't sound like how Adam from VH Partners usually writes" instinct when an invoice hits a shared AP queue.



Billing Account Update



Billing account update involves a request to change the bank account or payment routing for an existing vendor relationship, with the goal of redirecting future payments to an attacker-controlled account. Unlike a fraudulent invoice, which can blend into a queue of similar transactions, a billing account update asks the recipient to redirect where real money flows on an ongoing basis. Most organizations with any payment controls will pause on such a request, and threat actors know it. That friction explains why billing account updates carry the highest compromise rate of any VEC pretext at 26.5%.

Formal Controls, Higher-Effort Attacks

Government agencies have the highest billing account update rate of any industry at 40.8% of VEC—approximately 1.7x the sample average of 23.6%. The account compromise rate within those billing update attacks is also elevated at 41.2%, compared to the 26.5% sample average for that pretext. The government VEC sample is relatively small, meaning these figures are directionally suggestive rather than definitive, but the pattern is consistent with what the procurement environment would predict.

Vendor payment changes in government contexts typically require documentation, formal approval workflows, and audit trails. That procedural friction makes a convincing-sounding email from a lookalike domain a harder sell than it would be in a different environment, pushing attackers toward the higher-effort approach of compromising an actual vendor account before requesting a change to banking details.

The broader VEC data for government agencies reinforces the direction of this pattern. The overall account compromise rate across all VEC attack types targeting government agencies is 20.2%, more than double the sample VEC average of 8.95%, suggesting the shift toward higher-effort techniques extends beyond billing updates alone.


Why Legal Professionals See More Billing Fraud

Recipients in legal and compliance roles see billing account update fraud at a rate of 41%—the highest of any job category. Legal professionals manage trust accounts, escrow arrangements, wire transfers, and vendor payments across a wide range of external parties, including outside counsel, expert witnesses, court reporters, forensic investigators, and settlement administrators. A request to update payment routing is unremarkable in that context.

The reason legal professionals are particularly susceptible to this pretext is because of the nature of these relationships. Many are short-term and project-specific—e.g., an expert retained for a single case, a court reporter engaged for a deposition series, or a forensic consultant brought in for one investigation. The recipient may not have a long baseline of prior communications to compare against, which makes it harder to recognize when a billing update request breaks the pattern. There’s no established rhythm of “this is how this vendor usually communicates” to serve as a check.



Request for Quote (RFQ)



Request for quote (RFQ) fraud operates on a different logic than the other VEC pretexts. Invoice inquiry and billing account update both exploit an existing vendor relationship; RFQ fraud creates one from scratch.

The threat actor poses as a prospective customer soliciting pricing or proposals—a cold inbound inquiry that requires no prior context and no impersonation of a known party. The goal is either to establish a vendor relationship that enables later financial fraud or to extract pricing and financial information directly. That makes it uniquely effective against roles and functions where cold inbound inquiries are routine rather than suspicious.

When the Lure Looks Like a Lead

Sales and business development recipients see RFQ fraud at a rate of 40.5%—approximately 2.6x the sample average of 15.3%. That tracks with the role itself: responding to inbound quote requests is a core sales function. Unlike most VEC pretexts, which succeed by mimicking an existing relationship, RFQ fraud exploits the fact that no prior relationship is needed. A cold inquiry from an unfamiliar company asking for pricing on a product or service isn't suspicious in Sales—it's a lead. The entire function is built around engaging with unknown external parties.

The deal-closing orientation of the role amplifies the exposure. Sales teams operate under pipeline pressure and are rewarded for responsiveness. An unanswered quote request is a missed opportunity, and that urgency reduces the scrutiny that might otherwise catch a fraudulent request. The attacker doesn't need to manufacture urgency the way they would with a fake invoice or billing update; the target's own incentive structure provides it.

Payment Inquiry




Payment inquiry—direct requests for ad hoc payments that don't reference an existing invoice or vendor relationship—accounts for 19.8% of high-risk VEC campaigns, making it the third most common pretext after invoice fraud and billing account updates.

Unlike those two pretexts, payment inquiry doesn't produce notable variation across industries, org sizes, or job categories. That flatness is itself informative. Payment inquiry attacks are essentially the reconnaissance layer of VEC: a low-commitment probe to determine whether a target will engage before investing in a more elaborate approach. Threat actors taking this approach aren't looking for a specific type of organization or role; they're casting broadly to find whoever will respond. The absence of a targeting signal reflects the absence of targeting intent.



Regional Threat Profiles

VEC Pretext Mix by Region



Vendor fraud is not a monolithic threat. Attackers adapt their pretexts to regional business practices, resulting in a meaningfully different VEC landscape depending on the target organization's location.⁵

North America (NAM) and Asia-Pacific (APAC) share a broadly similar profile, with invoice fraud as the clear dominant tactic. Europe, the Middle East, and Africa (EMEA) diverges sharply, and the pattern suggests that regional business practices shape not just which pretexts threat actors favor, but how they structure entire fraud campaigns.

In NAM, invoice fraud accounts for 42% of VEC campaigns, followed by billing account update (24%), payment fraud (20%), and RFQ (14%). APAC tracks closely: invoice fraud leads at 40%, with payment fraud, billing account update, and RFQ all falling in a similar range to NAM's metrics. Both regions share a transactional procurement culture—less formal tendering, more direct invoicing, and a general orientation toward payment speed over process. In that environment, a fake invoice is the path of least resistance, and attackers take it.

While EMEA follows the same general distribution shape—one dominant pretext well ahead of the other three, which cluster within a few percentage points of each other—the lead pretext flips. RFQ fraud tops the list at 41%, nearly three times the NAM rate, while invoice fraud drops to 17%. It's the same level of concentration, just a different entry point—a sign that threat actors are deliberately matching their lead tactic to regional procurement norms.

⁵ Because the regional findings draw on materially different sample sizes, the non-NAM patterns should be read as directionally suggestive rather than definitive.



The EMEA sample is anchored by European organizations, and the regional pattern reflects that. European markets frequently require formal tendering processes, particularly for government and public-sector contracts. Cross-border complexity—multiple currencies, VAT systems, and regulatory frameworks—means structured quote processes are routine even in the private sector. RFQs are a standard part of vendor onboarding, and attackers exploit that familiarity. Stricter accounts payable controls and multi-party approval workflows in these markets may also render cold invoices a less reliable entry point, giving threat actors an additional reason to lead with procurement-stage pretexts that align with how business is already conducted.

The difference in dominant pretext also changes the shape of the threat for defenders. Invoice fraud—NAM's lead tactic—is fundamentally transactional: it arrives with a specific dollar amount, a payment deadline, and account details, all of which provide concrete signals for detection.

RFQ fraud operates differently. The initial contact contains no financial ask at all; it reads as a routine procurement inquiry, indistinguishable on its surface from a legitimate business inquiry. The financial payload comes later, after a relationship has been seeded and trust established. For organizations operating in RFQ-heavy environments, the challenge is recognizing vendor fraud at the relationship-building stage rather than at the payment stage—before the attack looks like an attack.



Real-World Example of Vendor Email Compromise

From: "S O" <s @ >
Sent: 9/18/2025 6:45:19 PM
To: "Accounts Payable" <accounts payable@ >
Cc: "G B" <g @ >
Subject: [Caution - External Email] Re: - Overdue Strategic Planning Services Invoice INV172573

Hi Accounting team

Following the request from W H below, I am forwarding INV172573 for processing as directed.

Kindly note that this invoice is already overdue, and we would appreciate it if it could be processed at your earliest convenience. Please confirm receipt.

S O

Accounting |

From: W H <w h @ >
Sent: Wednesday, September 17, 2025 04:31 PM
To: G B <g @ >
Cc: G B Accounting <s @ >
Subject: Overdue Strategic Planning Services Invoice - INV172573

I sincerely apologize for the delay in processing this payment; I had not seen your previous email until now. I have reviewed the invoice and everything is in order for payment.

To expedite the payment process and avoid any further delays please forward the invoice directly to our Accounting Department at accounts payable@ for processing

I want to acknowledge the exceptional quality of work your team provided during our strategic planning initiative in June. The comprehensive market analysis and strategic roadmap have already begun to influence our decision-making processes, and the executive workshops were particularly valuable for our leadership team.

Thank you for your patience and professionalism.

Regards

W H

From: G B <g @ >
Sent: Wednesday, September 17, 2025 10:31 AM
To: W H <w h @ >
Cc: G B Accounting <s @ >
Subject: Overdue Strategic Planning Services Invoice - INV172573

Hi W H,

I'm writing to follow up regarding the attached invoice issued for the recent strategic planning consultation services for Q2 provided to . As of today, the invoice remains outstanding and is now due.

Invoice Details:

- **Invoice Number:** INV172573
- **Invoice Date:** July 7th 2025
- **Amount Due:** \$49,130

We kindly ask that payment be processed at your earliest convenience. Please note that, as the invoice is overdue, it may be subject to a late payment fee in accordance with our payment terms.

Thank you again for the opportunity to support your team. We truly value our partnership and look forward to continuing to work together.

Best Regards

G B
 Chief Executive Officer

This email and any attachments are confidential and intended solely for the use of the named recipient(s). If you have received this email in error, please notify the sender immediately and permanently delete it from your system. Any unauthorized review, use, disclosure, copying, or distribution of this email is strictly prohibited. accepts no liability for any errors or omissions in the content of this email that may have occurred during transmission.

This attack begins the way many legitimate vendor interactions do: with a polite follow-up on an overdue invoice. The shared mailbox for the accounts payable department at a financial institution receives an email from S.O., an accounting contact at a consulting firm, forwarding an invoice for "strategic planning services." The message is professional, detailed, and aligned with a believable engagement between the bank and an external vendor.

What makes it compelling is the fabricated approval thread embedded below S.O.'s note. In the quoted conversation, the bank's own employee, W.H., apologizes for the payment delay, confirms that "everything is in order for payment," and instructs the vendor to forward the invoice directly to [accounts payable@\[redacted\].com](mailto:accounts payable@[redacted].com) for processing. Beneath that sits an earlier message from the vendor's CEO, G.B., complete with line-item detail, due date, and a warm acknowledgment of the bank's business. To a busy AP analyst, this appears to be a fully documented, stakeholder-approved obligation.

The attachments reinforce that illusion. The invoice is a polished consulting bill listing workshops, market analysis, and roadmap development, with ACH instructions to a Bank of America account in G.B.'s name. A matching W-9 provides taxpayer details and an electronic signature—everything the target needs to legitimize the vendor and push payment without additional verification.



Invoice
INV172573

Bill To

Project

TOTAL
\$49,130
Due Date: July 7, 2025

Terms	Issue Date	Due Date	Discount	Account Manager	Partner
Due On Receipt	July 7, 2025	July 7, 2025	\$0.00		

Item	Item Name	Quantity	Rate	Amount
L102-FD	Market and competitor analysis - Identification of key competitors, their strategies, and market positioning - Evaluation of competitor strengths, weaknesses, opportunities, and threats (SWOT) - Benchmarking of operational and financial performance against industry peers	1	\$21,130.00	\$21,130.00
ST-WL	Strategic planning workshops with leadership - Scenario planning and decision-making simulations for critical business initiatives	1	\$16,000.00	\$16,000.00
01-Setup	Setup Charge - - Including Initial setup and onboarding, executive team profiling - Establishment of communication protocols, reporting structures, and framework	1	\$2,000.00	\$2,000.00
TXT-RP	Development of strategic roadmap & action plan - Documentation of deliverables suitable for executive review, Board presentation, and cross-departmental coordination - Detailed Reporting	1	\$10,000.00	\$10,000.00

Subtotal

\$49,130.00

Amount Paid

\$0.00

Discount

\$0.00


Amount Due

\$49,130.00

*Invoices not paid within the due date may be subject to a late payment fee.

* Payment Instructions

ACH/Wire Information:
Bank Name:
Bank Address:
Account Name:
Account Number:
Routing number ACH/EFT:
Routing number DOM. WIRES:
SWIFT Code INTL WIRES:



Under the hood, however, this is pure impersonation. The entire “thread” is static HTML text inside a single email; there are no corresponding messages from W.H. or G.B. in the headers. The attacker operates a Microsoft 365 tenant on a domain they control, with DKIM passing and no restrictive SPF or DMARC policies. There are no authentication failures, and Microsoft assigns the message a low spam score.

The email contains no malicious links, payload-bearing attachments, or branding abuse—just a realistic vendor, a believable relationship, and banking details embedded in a convincing invoice. Together, the campaign’s elements perfectly mimic normal business paperwork, weaponize trust in internal approvers, and exploit routine invoice workflows to divert nearly \$50,000 to an attacker-controlled account.



Conclusion

- ■ The findings in this report converge on a single operational reality: modern email attacks are shaped by the institutions they target.

Phishing techniques calibrate to the defensive environment, with evasion methods that scale in sophistication alongside the security infrastructure they expect to face. BEC tactics reconfigure around organizational structure, swapping impersonation targets as the size and complexity of the workforce change which identities are credible to exploit. VEC pretexts align with regional procurement norms, reflecting an awareness of how business is conducted in the markets being targeted. In each case, the attacker's approach is less a static playbook and more a set of adaptive decisions driven by the operational context of the target.

That adaptability carries a specific implication for defense. When attacks are designed to blend into legitimate workflows, the signals that distinguish them from normal business communication are behavioral, not technical. A fraudulent invoice looks like a real invoice. A spoofed colleague request reads like a genuine one. A compromised internal account sends messages that are, by definition, indistinguishable from authentic internal email at the infrastructure level. Traditional security tools built around signatures, reputation lists, and static rules struggle with threats that are engineered to avoid triggering those mechanisms.

The data also challenges some widely held assumptions about where risk concentrates. VIP impersonation—the tactic most associated with BEC in the public imagination—accounts for a small fraction of internal impersonation attacks at large organizations, where employee impersonation and lateral attacks are far more prevalent. Vendor fraud, not internal social engineering, comprises the majority of BEC. And the most exposed mailboxes aren't necessarily the ones belonging to senior executives; shared functional inboxes like accounts payable sit at the intersection of high volume, low relationship context, and direct access to payment systems.

Defending against threats that mirror normal business requires understanding what normal business looks like—the communication patterns, approval workflows, vendor relationships, and identity structures that define each organization's operating environment. That understanding has to be specific, continuously updated, and applied at the point where a message is evaluated, not after the fact. The attacks documented in this report are designed to exploit the gap between how organizations actually operate and what their security tools are equipped to recognize.

Closing that gap requires AI that analyzes identity, context, and content to build behavioral baselines for every employee and vendor in an enterprise's cloud environment. That's what makes it possible to flag the moments when an attack tries to pass as business as usual—before an employee ever has the opportunity to engage.

Appendix

Methodology

The findings presented in this report are drawn from two complementary datasets covering attacks observed across Abnormal AI customers between July 1 and December 31, 2025. The primary dataset consists of 796,505 messages—a statistically representative 0.5% random sample of 159.4 million portal-visible attacks detected during the period—drawn from 4,669 customer accounts spanning three regions, 43 countries, and 21 industry categories. All statistics are derived from aggregated data; no individual customer, recipient, or account is identifiable in any published figure.

Simple random sampling was used to preserve attack type distributions and customer segment ratios, yielding a margin of error of $\pm 0.11\%$ at 95% confidence. A supplemental dataset covering vendor email compromise (VEC) was constructed as a census of VEC attack campaigns with verified categorical labels, analyzed at the campaign level (one representative message per campaign) to ensure statistical independence.

Analysis is limited to serious threats—phishing, business email compromise (BEC), and other malicious attacks—and excludes spam, graymail, and unwanted mail. Attack categories are assigned using a priority-based classification hierarchy in which phishing takes precedence, followed by BEC, then “other”. Lateral attacks and vendor fraud are tracked as independent attributes that may apply across any primary category.

Demographic dimensions presented in this report—including industry, geography, organization size, and recipient job function—reflect the composition of Abnormal AI’s customer base and should not be interpreted as evidence of attacker targeting preferences. Because the sample is drawn from a defined set of customers rather than a representative cross-section of all organizations globally, the relative volume of attacks observed within any given industry, region, or role is a function of both attacker behavior and customer mix.

Organization Size Breakdown

- **Small:** 0 - 1,000 mailboxes
- **Mid-Market:** 1,001 - 5,000 mailboxes
- **Large Mid-Market:** 5,001 - 10,000 mailboxes
- **Enterprise:** 10,001 - 50,000 mailboxes
- **Large Enterprise:** 50,001+ mailboxes

Where demographic dimensions appear in this report, they are used as analytical slicing variables to compare attack characteristics within segments (e.g., how attack types or techniques differ across industries), not to rank or compare how frequently segments are targeted in absolute terms. Content-level statistics—including link analysis, phone number detection, and vocabulary signals—are derived from real-time processing logs available for 99.5% of sampled messages; the remaining 0.5% of messages, processed via alternative detection paths, are excluded from content-level figures but are included in all high-level attack classification statistics.



Definitions

Phishing: Attacks that attempt to steal credentials through malicious links, QR codes, or attachments, typically by directing recipients to fraudulent login pages designed to harvest usernames and passwords.

Business email compromise (BEC): Social engineering attacks that impersonate trusted parties—e.g., employees, executives, internal departments, partners, or vendors—to manipulate recipients into taking a harmful action. Encompasses all impersonation-based social engineering that does not meet the criteria for phishing.

Vendor email compromise (VEC): A subset of BEC in which attackers impersonate a vendor or business partner to pursue financial fraud. Distinguished from other BEC attacks by its focus on vendor relationships and financially-motivated outcomes. Categorized by four attack goals: Billing Account Update, Invoice Fraud, Payment Fraud, and Request for Quotation. Vendor impersonation targeting non-financial outcomes is classified as BEC, not VEC.

- **High-risk VEC:** VEC campaigns exhibiting significant financial risk or sophisticated attack techniques. These campaigns undergo additional analysis with verified categorical labels, enabling detailed breakdowns by attack goal, impersonation technique, and targeting pattern.
- **Standard VEC:** VEC campaigns involving vendor or partner impersonation that fall below the High-Risk threshold.
- **Vendor impersonation:** A VEC technique in which the attacker creates a lookalike email address or spoofs the sender to appear to be a known vendor. The vendor's actual account is not compromised.
- **Vendor compromise:** A VEC technique in which the attacker has taken over a vendor's actual email account, and the malicious message genuinely originates from the vendor's domain.

Lateral attack: An attack sent from a compromised internal account within the target organization, exploiting existing trust between colleagues. Can occur across any attack category.

Other: Attacks that do not meet the criteria for phishing or BEC, including malware, extortion, ransomware, and social engineering that does not impersonate a known or trusted party.

• • × × ⊙ ⊙ ⊙ ⊙

Acknowledgments

This report was produced by the Abnormal AI threat intelligence and marketing teams. The authors would like to acknowledge the following individuals for their contributions to the research, analysis, and writing of this report: Callie Baron, Elizabeth Swantek, Vincent Aiello, Erin Ludert, Steven Kendall, Andrew Rud, Izzy Tantillo, and Piotr Wojtyla.





About Abnormal AI

Abnormal AI is the leading AI-native human behavior security platform, leveraging machine learning to stop sophisticated attacks and detect compromised accounts across email and connected applications. Our anomaly detection engine leverages identity and context to analyze normal behavior and assess the risk of every cloud email event—detecting and stopping sophisticated, socially-engineered attacks that target your organization's most valuable cybersecurity asset: your people.

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

Interested in Stopping Modern Email Attacks?

[Request a Demo >](#)