




Abnormal

The Essential Guide to Business Email Compromise

How BEC Became a \$3 Billion Problem,
and What It Takes To Stop It

► What Business Email Compromise Looks Like Today

Cybercriminals steal billions without sending a single link or attachment.

At 61%, vendor impersonation is the biggest driver of business email compromise, bigger than CEO fraud.

In 2025, AI-enabled fraud grew so prevalent the FBI gave it its own category for the first time.

Why is BEC Still Winning?

Organizations have spent years hardening their defenses against email attacks. Secure email gateways scan for malicious links and attachments. Employees complete awareness training while finance teams add identity verification steps. The SOC refines detection rules. And yet, business email compromise remains one of the most financially damaging forms of cybercrime in the world.

In 2025, the [FBI's Internet Crime Complaint Center](#) received 24,768 BEC complaints representing more than \$3 billion in reported losses, roughly \$123,000 per incident. Perhaps even more alarming than the amount of losses is the rate of exposure: even organizations with fewer than 1,000 employees face a [70% weekly probability](#) of receiving at least one BEC attack.

This is the paradox at the heart of modern BEC. The attacks causing the greatest financial damage often contain none of the indicators legacy security tools were built to detect.

Traditional Email Inspection

Sender	Pass	✓
Authentication	Pass	✓
Links	Pass	✓
Attachments	Pass	✓
Domain Reputation	Pass	✓
Result	Deliver to Inbox	✓

What the Legacy Inspection Never Evaluated:

- Is this request expected?
 - Has this executive ever emailed this employee?
- Has this vendor changed bank details before?
 - Is this consistent with normal behavior?

Even a single successful BEC attack can redirect millions of dollars, compromise a critical vendor relationship, and expose sensitive financial data. That's why BEC remains a board-level concern in a crowded threat landscape.

So why does BEC keep succeeding, at this scale, despite everything organizations have thrown at it?



How BEC Evolves

Business email compromise is often treated as a single, stable attack type. It's neither. The vector has evolved through four layers of sophistication, each one building on the last rather than replacing it. Modern BEC can draw on all four at once, combining executive authority, vendor credibility, the inherited trust of a real account, and AI-generated personalization in whatever mix the attacker needs.

Layer 1 | Executive Impersonation

The foundation, and where BEC started. Attackers pose as CEOs, CFOs, or other senior leaders and lean on urgency and authority to pressure an employee into wiring funds or sharing sensitive information. The premise is simple: people are more likely to comply with a request from someone senior. It is the oldest form of BEC, and it still works today.

In 2015, a single email impersonating the CEO of FACC, an Austrian aerospace manufacturer, convinced a finance employee to wire €42 million toward what was described as a confidential acquisition. By the time the fraud surfaced, most of the money was gone.

Layer 2 | Vendor Fraud

Attackers soon found that targeting vendor relationships scaled better than impersonating internal executives. Instead of manufacturing urgency, they trade on credibility: a payment update, a revised invoice, new banking details, requests that look routine because they are the kind of transaction an accounts-payable team processes every day.

The approach has proven durable. Vendor fraud now accounts for 61% of all BEC, and 81% of attacks targeting shared mailboxes like accounts-payable@ and invoices@ are vendor-related. The average loss from a single vendor email compromise is \$100,000.

In December 2022, ETS Contracting received a payment-instruction email it believed came from the owner of a vendor it worked with, directing it to wire \$240,926 to a Chase account. The email was fraudulent, sent from a lookalike domain that changed a single letter, and the account belonged to the scheme's perpetrator rather than the real vendor. ETS authorized the payment, and the funds posted the next day before the deception was discovered.



How BEC Evolves (Continued)

Layer 3 | Compromised Accounts and Conversation Hijacking

The next layer goes further. Rather than impersonating a trusted contact from the outside, attackers operate from inside a real conversation, either by compromising a legitimate account or by inserting a lookalike into an active thread. They read the history, learn the workflow, and wait for the natural moment to redirect a payment. The sender looks right. The thread is real. The request fits.

Beginning around August 2022 and prosecuted through 2026, a transnational criminal organization ran a business email compromise operation that stole more than \$38 million from individuals and businesses in the US and abroad. The group gained access to victims' email accounts and used that access to insert fraudulent payment instructions into legitimate financial correspondence, redirecting wires to accounts they controlled. Four defendants were sentenced for their roles in April 2026 after pleading guilty to wire-fraud and money-laundering conspiracies.

The shift that account takeover represents is at the heart of modern BEC. Attackers don't have to manufacture trust. They inherit it.

Layer 4 | AI Acceleration

Generative AI compresses every stage of the attack. Bad actors now craft convincing executive impersonations, mirror a specific person's writing style, localize messages, and personalize at scale, all with widely available tools.

The scale of the shift is significant enough that the FBI began tracking AI-enabled crime as its own category for the first time in its 2025 report, logging more than 22,000 complaints and nearly \$900 million in associated losses. The report describes the exact capability that makes this layer so dangerous: AI drafting emails that match a CEO's writing style closely enough to pass for the real person.

What once required real time and skill now takes neither. The level of personalization and apparent legitimacy that used to mark the most resourceful attackers is now available to anyone.

Every layer exploits the same thing: behavior. The tools and the tactics change, but the goal remains consistent. Convince someone to act on familiarity, trust, and the ordinary expectations of their workday.

Explore the 2026 Attack Landscape Report for critical threat intelligence on attacks targeting your organization.


Get the Report >



Why Traditional Defenses Struggle

There's a reason BEC keeps succeeding even where organizations have layered on multiple email defenses. Traditional email security was built to answer one question: does this message look malicious? Modern BEC demands a different one: is this behavior normal?

A traditional gateway can evaluate domains, scan links, inspect attachments, and check messages against known threat indicators. What it cannot do is understand the relationship between a sender and a recipient, judge whether a request fits an established pattern of communication, or recognize that a legitimate account has been compromised and is now being operated by someone else. When the account is real, the thread is real, and the message carries no link or attachment, there is nothing for a content scanner to catch.



Technical Check

SPF	Pass	✓
DKIM	Pass	✓
DMARC	Pass	✓
Sender Domain	Legitimate (proliasystems.com)	✓
Links	None	✓
Attachment	PDF, No Malware	✓

Behavioral Signals

Origin	First-ever send from a new location (1,300+ prior emails, none from here)	⚠
Reply-to	Silently changed to a lookalike domain	⚠
Request	First banking-detail change in the history of this relationship	⚠
Urgency	Hard same-day deadline, atypical for this sender	⚠

Every technical check passes. Every behavioral signal fails. A gateway sees the left column. Only a behavioral model sees the context that makes the request suspicious.



Why Traditional Defenses Struggle (Continued)

Modern BEC compounds the problem by aiming at the softest structural targets. Shared mailboxes such as accounts-payable@, invoices@, and purchasing@ are publicly discoverable, handled by several people, process high volumes of payments, and are the least likely to hold deep relationship history on any given sender. They are, in other words, built for exactly the kind of routine vendor request an attacker wants to imitate.

Security awareness training has improved, and robust verification procedures have helped. But neither was built for an attack with no artifact to spot and no anomaly a person would think to question, one that arrives looking exactly like the routine business it imitates. What gives it away isn't in the message at all. It's in the behavior behind it.

The hardest BEC attacks don't look like attacks. They look like another message in an ordinary business conversation.



What Effective BEC Defense Looks Like

The message on the previous page passed every technical check and failed every behavioral one. That gap is the scope of the problem, and closing it means answering a different set of questions than a legacy SEG was built to ask; questions about identity, relationship, and intent, answered in real time, across every message in the organization.

Who is actually communicating?

Authentication confirms that a message passed technical validation. It does not confirm identity. Knowing who a sender really is takes visibility into communication history, established relationships, and behavioral patterns, not just whether SPF and DKIM passed. The compromised vendor account in our example cleared every authentication check.

Is this relationship established?

Has this coworker, vendor, or IT contact ever communicated with this recipient before, and if so, how often and in what context? A first-contact payment request carries a fundamentally different risk profile than a message inside a long-standing relationship.

Is this request expected?

Does the business context hold up? A payment update arriving mid-conversation is one thing. An unsolicited banking change from a sender with no prior history is another. Context can't be judged from a single message read in isolation.

Does the behavior match history?

Every sender has patterns: how often they write, what they typically ask for, their style, their working hours. A first-ever banking change wrapped in a same-day deadline is exactly the kind of deviation worth acting on, even when every technical indicator looks clean.

Should this be trusted?

One sole signal cannot answer this. The answer emerges from synthesizing identity, relationship, context, and behavior together.

True behavioral AI answers these questions at scale. It continuously builds a model of every identity in an organization's communication network, maps how relationships behave over time, and flags the anomalies, including the ones hiding inside legitimate, authenticated accounts. It understands not just what a message contains, but who sent it, why, and whether the behavior behind it fits what the organization knows to be normal.



BEC Defense Evaluation Checklist

- ☐ Does it verify identity, not just authentication?
- ☐ Does it understand relationships, not just message content?
- ☐ Can it assess whether a request fits the business context?
- ☐ Can it detect behavioral anomalies in legitimate, authenticated accounts?
- ☐ Can it apply differentiated risk thresholds to high-volume shared mailboxes?
- ☐ Does it continuously adapt, or does it require manual rule-writing and tuning?

For deeper evaluation of your current defenses, download The Essential Guide to Cloud Email Security.

[Get the Guide >](#)



The New Reality of BEC

BEC has become one of the most persistent and costly categories of cybercrime in the world, and the trajectory points up, not down. Reported losses rose again in 2025, despite BEC being a widely publicized threat to organizations of all sizes.

The four layers trace how it got here. What began as external executive impersonation became a threat that operates from inside real accounts, rides real conversations, and generates flawless, personalized messages on demand. Each layer added a new source of credibility without retiring the last. AI removed the last constraint: the time and skill it once took to make an attack convincing. The messages landing in inboxes will only grow harder to separate from legitimate business.

The five truths at the start of this guide describe where that leaves us. BEC is a financially damaging, continuously evolving threat that exploits trust rather than technology, targets people rather than systems, and succeeds precisely because it looks like the ordinary business it imitates. No identity, relationship, or channel is safe from it by default.

Defenders must adapt to the acceleration of BEC via AI by adopting defensive AI fueled by behavioral intelligence. Detecting malicious content was the right goal when content carried the attack payloads. Modern BEC carries none. Stopping it means recognizing when trusted, authenticated, clean communication is behaving in a way that doesn't fit the organization's "normal" baseline, and acting before a human ever has to make the call.

That is the shift the last decade has forced: from inspecting messages to understanding behavior.



►► 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.

Uncover your organization's real BEC exposure.

[Schedule a Demo ►](#)

See how Abnormal stops BEC with behavioral AI.

[Take a Product Tour ►](#)