Abnormal

WHITE PAPER

ABX: Abnormal Behavior Technology

OCTOBER 2025

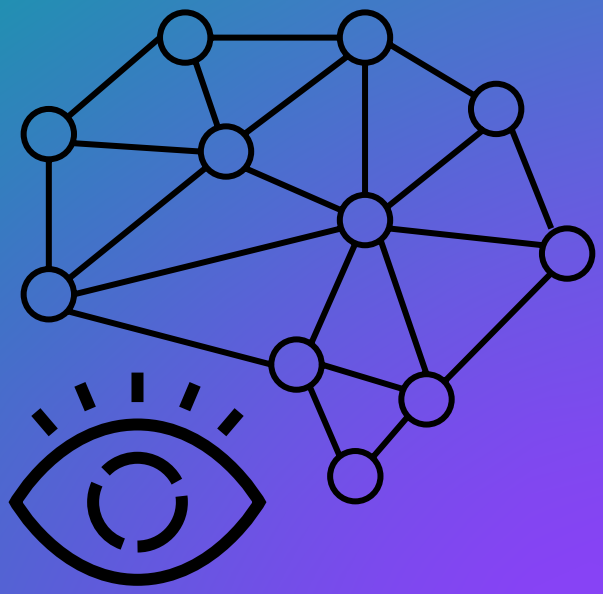

Executive Summary

$55.5B

in exposed global losses from BEC between 2013 and 2023

IC3 2024 Internet Crime Report

$2.77B

lost to BEC in 2024 alone

IC3 2024 Internet Crime Report

Email remains the foundation of business communications and the primary entry point for cyberattacks. The rise of generative AI has revolutionized the cyber threat landscape, giving adversaries the ability to craft highly convincing phishing messages, scale business email compromise campaigns, and impersonate executives and vendors with unprecedented realism. Despite billions spent on secure email gateways, security awareness, and layered legacy defenses, organizations continue to suffer staggering losses.

Business Email Compromise (BEC) stands among the most financially devastating cyberthreats. The FBI's Internet Crime Complaint Center (IC3) reports **$55.5 billion** in combined exposed losses from BEC globally between 2013 and 2023. In 2024 alone, BEC accounted for **$2.77 billion** in losses, second only to investment fraud.

To bypass traditional defenses, attackers increasingly rely on social engineering enhanced by generative AI. These techniques allow them to mimic trusted identities, tailor messages with context-specific detail, and evade detection by secure email gateways. Others exploit connected applications and third-party integrations to gain direct access to the email tenant. As AI accelerates the speed and realism of these campaigns, protecting the enterprise's most critical communication channel has become a strategic imperative.

Defending against AI-powered adversaries requires a different approach—one that looks beyond static indicators and payloads to understand behavior itself. Abnormal ingests thousands of internal and external signals through an API integration to establish a living model of how identities and applications normally operate. By grounding detection in this behavioral baseline, the platform exposes the minute anomalies that reveal social engineering, account takeover, or malicious third-party access.

Abnormal evaluates every event through three layers of behavioral intelligence. It becomes identity aware by building detailed profiles of employees, vendors, and applications from directories, sign-in patterns, and communication histories. It is context aware, mapping the relationships among those identities and analyzing the tone, cadence, and frequency of their interactions. And it is risk aware, applying advanced natural language models and deep content analysis to detect suspicious intent, dangerous URLs, or hidden payloads. Taken together, these dimensions enable high-confidence detection of the most sophisticated social engineering attacks.

Unlike other solutions, Abnormal delivers this protection with no tuning and no disruption to mail flow. The API-native architecture integrates directly with Microsoft 365 and Google Workspace, deploying in minutes and continuously adapting as behaviors evolve. Security teams gain immediate visibility and precise detection, without the maintenance burden that often weighs down legacy defenses.

Table of Contents

The Problem With Email Security

Socially-engineered attacks such as business email compromise evade traditional email security because they lack the common threat signals that trigger a detection in most security solutions. These attacks do not have attachments carrying malware nor do they contain links leading to malicious websites. The content of the email is generally simple, and the attacks are typically customized for each individual target.

BEC attacks are nearly always hand-crafted and incorporate heavy elements of social engineering. While they represent only a narrow slice of overall attack volume, their precision makes them disproportionately costly and among the most financially devastating email threats.

Attackers are also expanding beyond BEC, exploiting misconfigurations, compromised accounts, and third-party integrations. Effective cloud email security must address these platform-level entry points to prevent account takeover and the internal attacks that follow.

Defending against these evolving threats requires a platform that understands identity and behavior to stop social engineering threats powered by AI and the next wave of never-before-seen attacks exploiting the human vulnerability.

Abnormal Behavioral AI applies behavioral intelligence to cloud email by ingesting thousands of internal and external signals. It builds behavioral baselines across identities and relationships, then analyzes each event through the lenses of identity, context, and risk. Using an ensemble of advanced models—including natural language processing and behavioral anomaly detection—Abnormal delivers high-confidence, explainable verdicts that block even never-before-seen attacks.

The Modern Email Attack Framework

Abnormal AI has developed the following framework to provide insight into how ABX identifies and addresses socially-engineered email attacks.

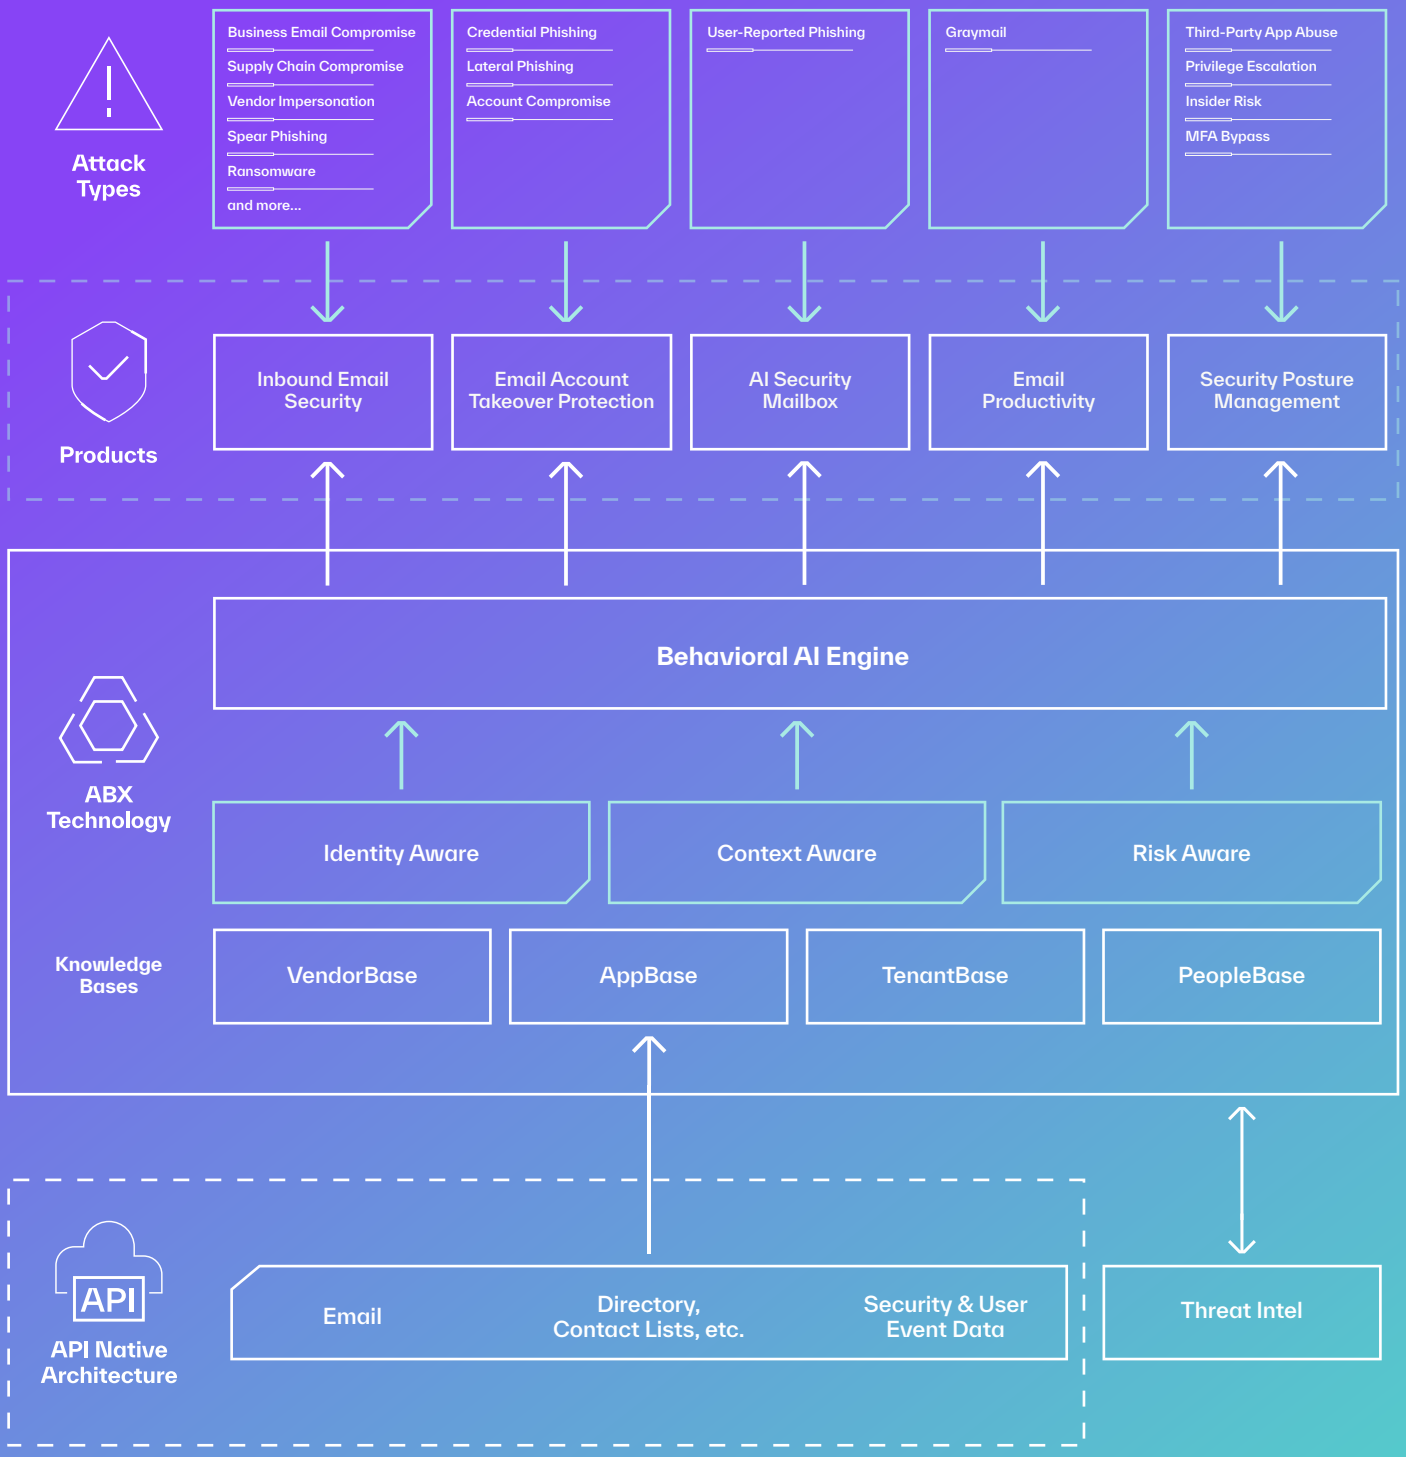

Attacks That Evade Legacy Defenses

Executive Impersonation

Pretext
Internal Employee

Approach
Impersonation

Delivery
No Payload

Executive impersonation is a hallmark of BEC, and its effectiveness lies in its ability to exploit authority and urgency. Attackers no longer rely on clumsy spelling errors or obscure domains—today, they use generative AI to craft polished messages that appear to come directly from senior leadership.

These emails often originate from widely trusted services such as Gmail or M365, making them indistinguishable at the domain level and impossible to block without disrupting legitimate communication.

Traditional rule-based approaches have proven brittle and unsustainable. Modern impersonation detection requires baselining a wide range of identity and behavioral signals: login location, device and IP reputation, typical communication hours, and writing style. When evaluated together with the tone, topic, and context of the message, these signals expose subtle anomalies that reveal even the most convincing impersonation attempts.

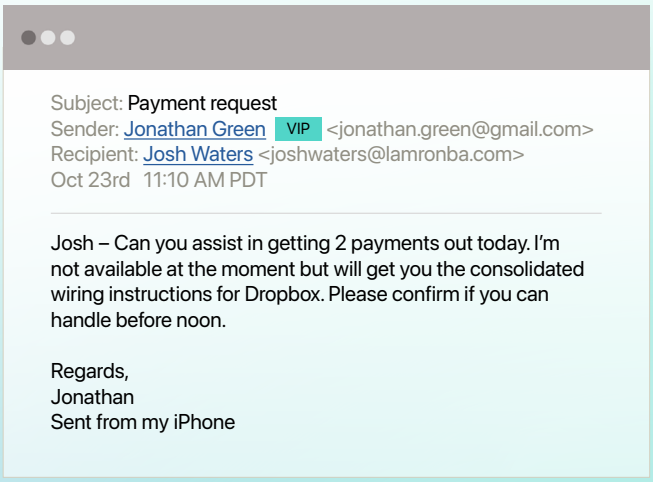

Subject: Payment request
Sender: Jonathan Green VIP <jonathan.green@gmail.com>
Recipient: Josh Waters <joshwaters@lamronba.com>
Oct 23rd 11:10 AM PDT

Josh – Can you assist in getting 2 payments out today. I'm not available at the moment but will get you the consolidated wiring instructions for Dropbox. Please confirm if you can handle before noon.

Regards,
Jonathan
Sent from my iPhone

Vendor Email Compromise

Pretext
External Partner

Approach
Compromised Account

Delivery
No Payload

Vendor email compromise is especially dangerous because it exploits trusted relationships. Once an attacker gains access to a vendor's account, their messages come from legitimate domains and may even appear within ongoing threads. This makes the messages indistinguishable from authentic correspondence. The appearance of credibility allows attackers to introduce fraudulent requests, often disguised as updated invoices or new banking details, that slip past traditional controls.

Detecting these attacks requires context awareness of normal vendor-customer relationships and risk awareness of expected content patterns. Security platforms must learn how vendors typically communicate—what topics they cover, how often they interact, and which contacts are usually involved—and flag anomalies such as unexpected financial requests or altered account information. Only by modeling this baseline of behavior can organizations reliably expose vendor compromise before it leads to financial loss.

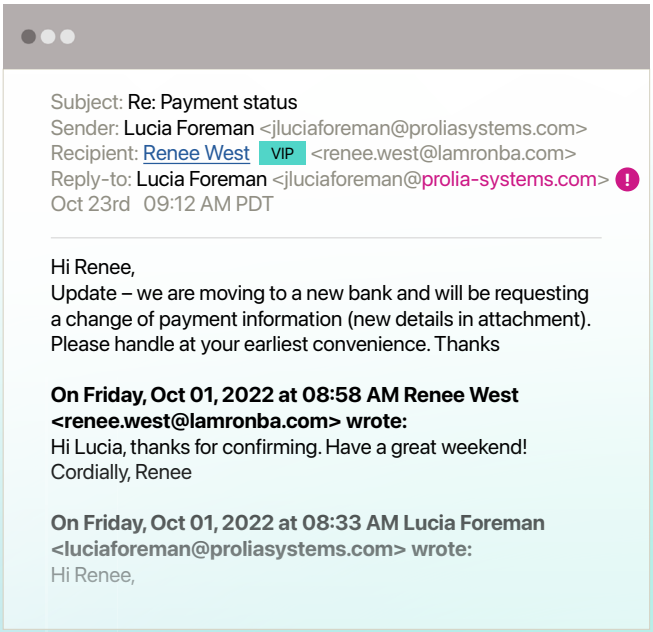

Employee Compromise

Pretext
Internal Employee

Approach
Compromised Account

Delivery
No Payload

Like compromised vendor accounts, attacks from internally compromised accounts are also extremely difficult to identify. The emails come from trusted employees, may reference legitimate business information, and may bypass security tools that do not scan east-west mailflow. Once an attacker has gained access to an internal email account, he can use it to uncover information, move through connected applications, or send additional attacks to employees and customers, often for long periods of time without detection.

Subject: Orion Limited Invoice
Sender: Renee West VIP <renee.west@lamronba.com>
Recipient: Josh Waters <joshwaters@lamronba.com>
Oct 23rd 02:46 PM PDT

Zachary and Josh,
Please review the attached – I have approved this wire transfer and it should be prepared for immediate release.

Thanks,
Renee West
Treasurer
www.lamronba.com

Orion Limited
2659 Cottrill Lane
Saint Louis MO 63146

INVOICE
Inv201822934

Bill To:
Lamronba Corporation
123 Main Street
San Francisco, California 98100
Attention: Zachery Newton

Date: Nov 20, 2018
Balance Due: $1,120,500.00

Item	Quantity	Rate	Amount
Y2018 Invoice to date	1	$1,120,500.00	$1,120,500.00

Credential Phishing

Pretext
Brand

Approach
Impersonation

Delivery
Link to Credential Phishing Website

Most credential phishing attempts use impersonation of a known brand, such as Microsoft, Amazon, LinkedIn, Google, or another large organization that the recipient is likely to recognize. And once the recipient has entered their credentials, the attacker can use them to gain access to the account and all associated information.

While some email security solutions may detect these attacks, particularly if they use high entropy or previously seen URLs, these attacks are difficult to reliably catch without the risk awareness and known-normal identity baselines. The fact that credential phishing sites typically do not contain malware makes typical sandboxing approaches ineffective.

Subject: Microsoft365 password expiry notice!!!
Sender: Acme Microsoft Support <microsoft@acme.com>
Recipient: Adam Smith <Asmith@acme.com>
Nov 24th 05:30 PM PST

Password expiry notice!!!

User name: asmith@acme.com

Here's what to do next:

- Click the link below.
- Use the button below to re-confirm and continue with the same password.

Re-activate Password

If this issue isn't resolved, your subscription and any data you may have stored in it will be permanently deleted on 31 November 2022.

Sincerely,
The Acme Microsoft Support Team

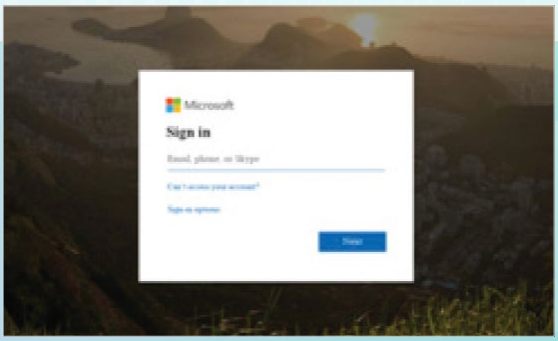

MFA Bypass

Pretext
Internal Employee

Approach
Exploiting Conditional Access Misconfiguration; Brute Force

Delivery
Legacy Authentication Exploitation; Push Notification Spam

Attackers don't always rely on inbound email—they often target accounts directly. When multi-factor authentication (MFA) is enabled, adversaries look for ways to bypass it. One common tactic is exploiting legacy authentication, downgrading to older mail clients or protocols that don't enforce MFA, and using stolen credentials to log in without additional checks. When legacy authentication is blocked, attackers frequently turn to MFA fatigue attacks, bombarding users with push notifications until they approve one out of frustration or confusion.

For defenders, visibility into MFA bypass attempts is critical. Without it, attackers can maintain undetected access to cloud email accounts for extended periods, escalating privileges or launching internal attacks.

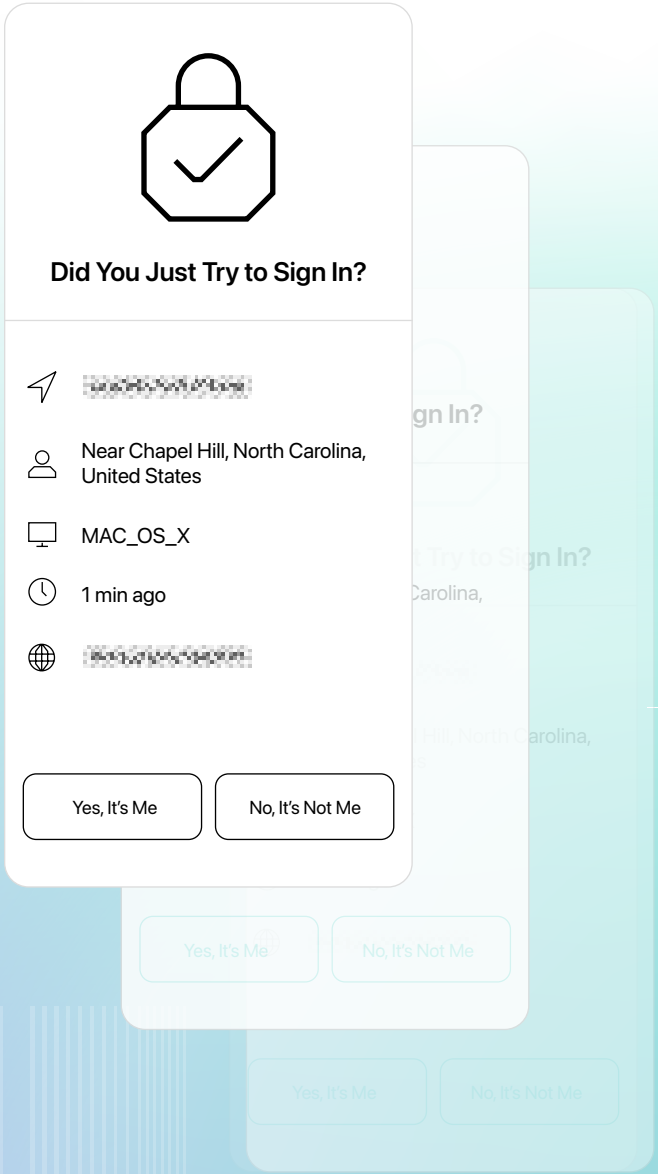

Malicious or Over-Permissioned Third-Party Applications

Pretext
Internal Employee

Approach
Updating App Permissions

Delivery
API Integration

Many third-party application attacks begin with a phishing email, but instead of stealing credentials, the attacker persuades the target to authorize a new application. Once approved, these apps often request extensive read/write permissions, giving adversaries deep access to calendars, mailboxes, and sensitive data. In other cases, attackers compromise legitimate tools—such as calendar add-ons or spellcheck services—and weaponize that trusted access to move inside the email platform.

Detecting these attacks requires visibility into app permissions and changes over time. Security teams must be able to surface newly authorized or over-permissioned applications and evaluate whether that access aligns with normal business needs before attackers can exploit the gap.

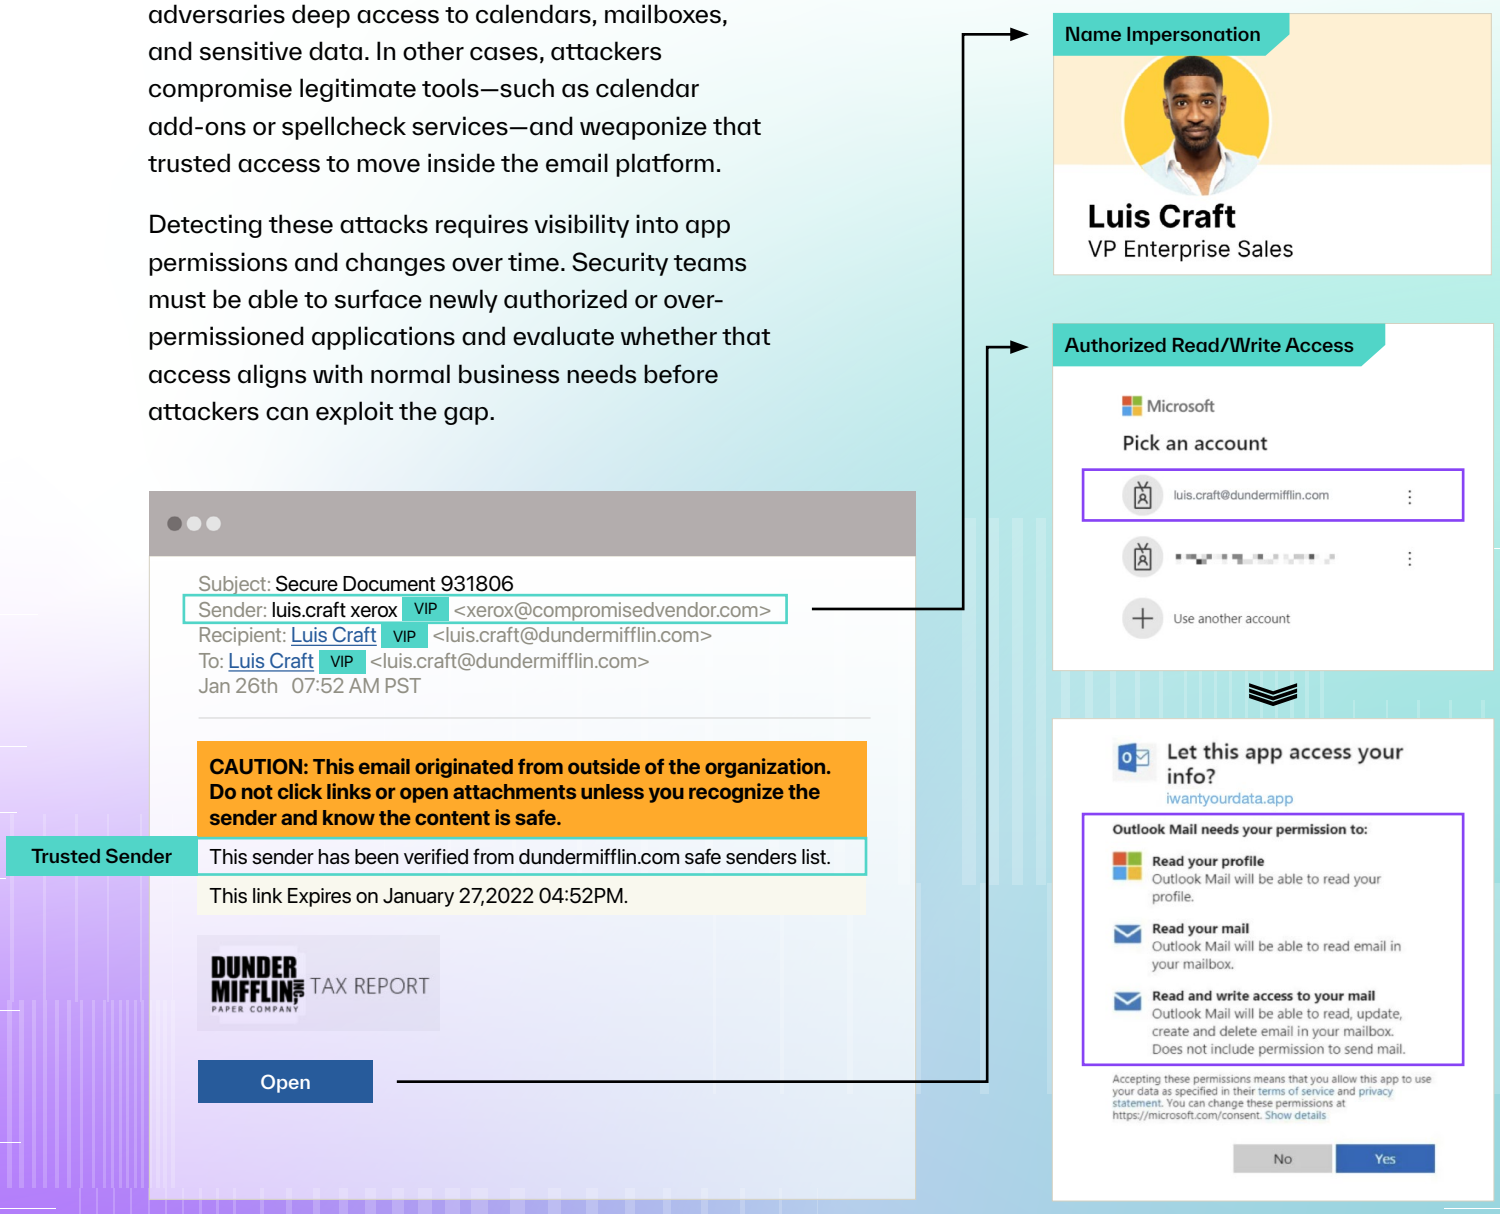

The Role of Identity in Stopping Attacks

In the wake of successful attacks, investigations from security teams extend past the email into accounts, identities, and applications. Examples of investigative activities may include:

- Identifying the sender and the target they impersonated
- Verifying executive personal accounts to confirm impersonation
- Confirming vendor bank account details when attackers request suspicious changes
- Reviewing logs of internal accounts for signs of compromise
- Reviewing MFA logs to confirm bypass methods
- Evaluating third-party app permissions to ensure they align with legitimate business needs

Many traditional email security solutions do not perform these activities when attempting to block socially engineered or platform-level attacks.

In contrast, ABX learns from each customer environment, uniquely leveraging a broad set of organization-specific data to protect the enterprise. By doing so, Abnormal can detect and stop the attacks that other solutions miss.

Inside Abnormal Behavioral AI

Abnormal Behavior Technology (ABX) powers Cloud Email Security by ingesting thousands of signals from your Microsoft 365 and Google Workspace environments. ABX profiles every user, vendor, and customer to establish behavioral baselines and context.

The platform uses ensemble AI to analyze identity, relationships, and risk across all email activity—enabling detection of the most advanced inbound threats, including business email compromise and account takeover, without manual tuning or frequent configuration changes.

ABX ingests and analyzes the rich data from dozens of data sources to profile communications across three distinct categories to provide:

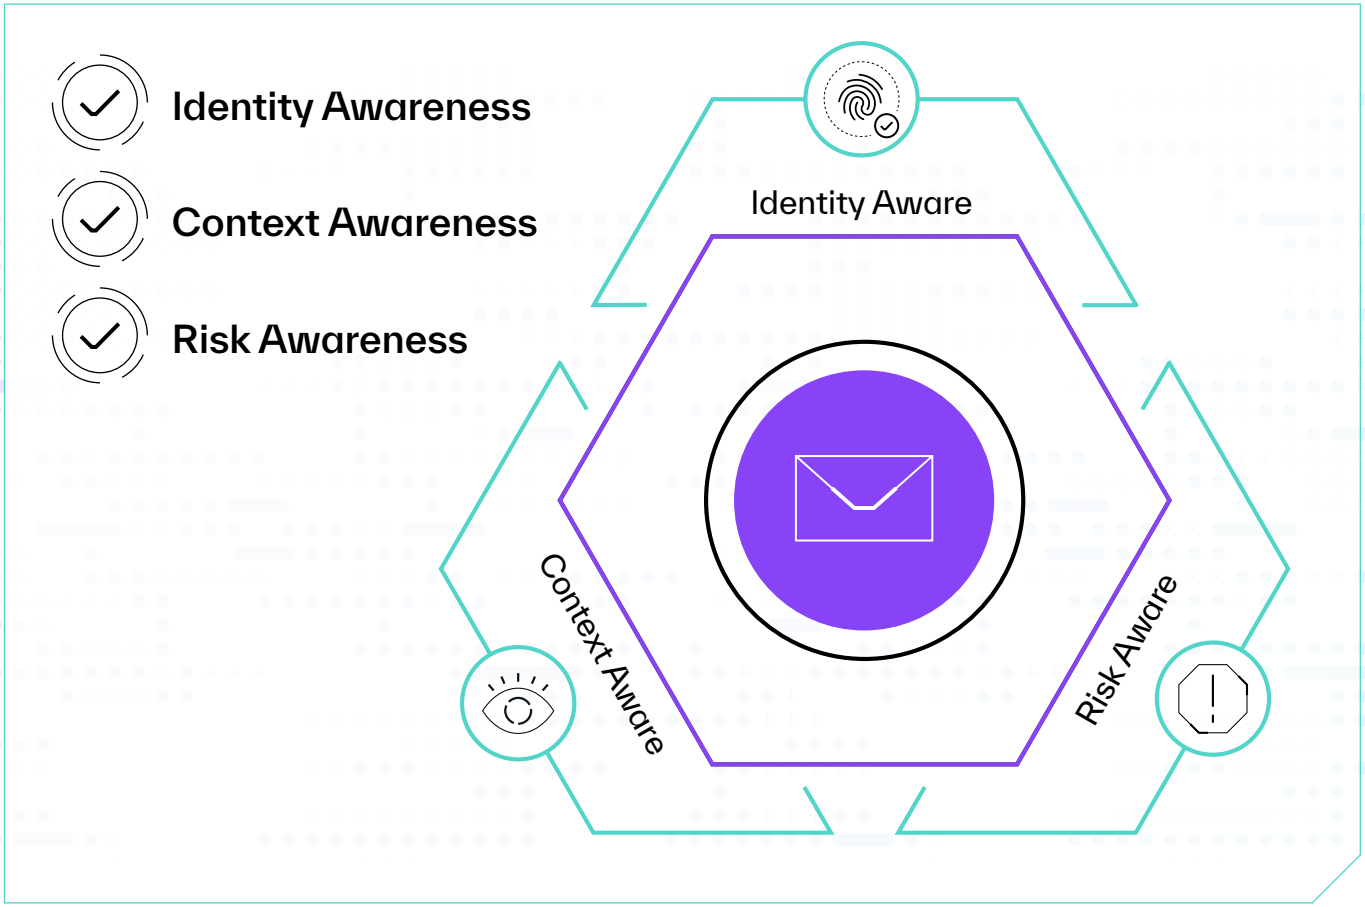

The results of the analysis across these three areas are then consolidated by an ensemble of machine learning algorithms to ensure a high-confidence verdict of whether a threat exists in the email platform—minimizing the false positives that plague traditional machine learning algorithms.

INSIDE ABNORMAL BEHAVIORAL AI \\ **01**

Abnormal Identity Awareness

Abnormal Identity Awareness is a stateful model that ingests thousands of internal and external identities. For employees, ABX takes inputs from the directory, analyzes user events, and analyzes email communications, resulting in models for each employee. The attributes for each internal identity include:

Employee Identity Model		
Name	Email	Role
Personal Email	Location	Sign-In Locations
Manager	Manager Location	Department
VIP Status	Office Address	Phone Number
Term at Company	Browsers Used	Devices Used
Usual Login Time	Mail Filter Configuration	Client Application Used
Mailing Address		

To create models for external entities, ABX evaluates the email communications in detail to extract identity attributes.

Vendor Identity Model		
Vendor Name	Email Used for Communication	Key Vendor Contacts
Key Internal Contacts	Mailing Address	Verified Email FQDN
Phone	Invoicing Software	Invoicing Cadence
Key Vendor Contacts	Key Internal Contacts	Bank Information / Accounts
Invoicing Language	Last Contacted	Phone
Years of Relationship		

Customer Identity Model		
Customer Name	Customer Emails	Key Customer Contacts
Key Internal Contacts	Mailing Address	Verified Email FQDN
Phone	Invoice Frequency	Last Contacted
Years of Relationship	Communication Cadence	

To create models to detect anomalous activity from third-party apps, ABX ingests and monitors the identities of each application to secure the organization's posture from malicious third-party apps.

Third-Party Application Model		
Installed Application	Application Name	Release Date
Version	Read Access	Write Access
Number of Applications		

To create models that monitor changes and configuration risks that exist in the tenant, ABX ingests and baselines behaviors of the identities in each tenant.

Tenant Model	
Number of Tenants	Name of Tenant/Platform
Permissions in Tenant	New Users Added
Date of New Users Added	

ABX ingests thousands of signals from both internal and external entities to build out baselines of the known normal behaviors of those identities, making it identity-aware.

INSIDE ABNORMAL BEHAVIORAL AI \\ **02**

Abnormal Context Awareness

After establishing identity baselines, ABX profiles communication patterns within your email tenant, analyzing the frequency, topic, sentiment, and tone of every email to create continuously updated context graphs. This approach identifies out-of-character communication, detects rare relationship paths, and recognizes anomalous message topics that could signal targeted attacks such as executive or vendor impersonation.

Ongoing enhancements in behavioral modeling and anomaly detection allow Abnormal to surface nuanced threats that may bypass traditional rule- or policy-based security tools.

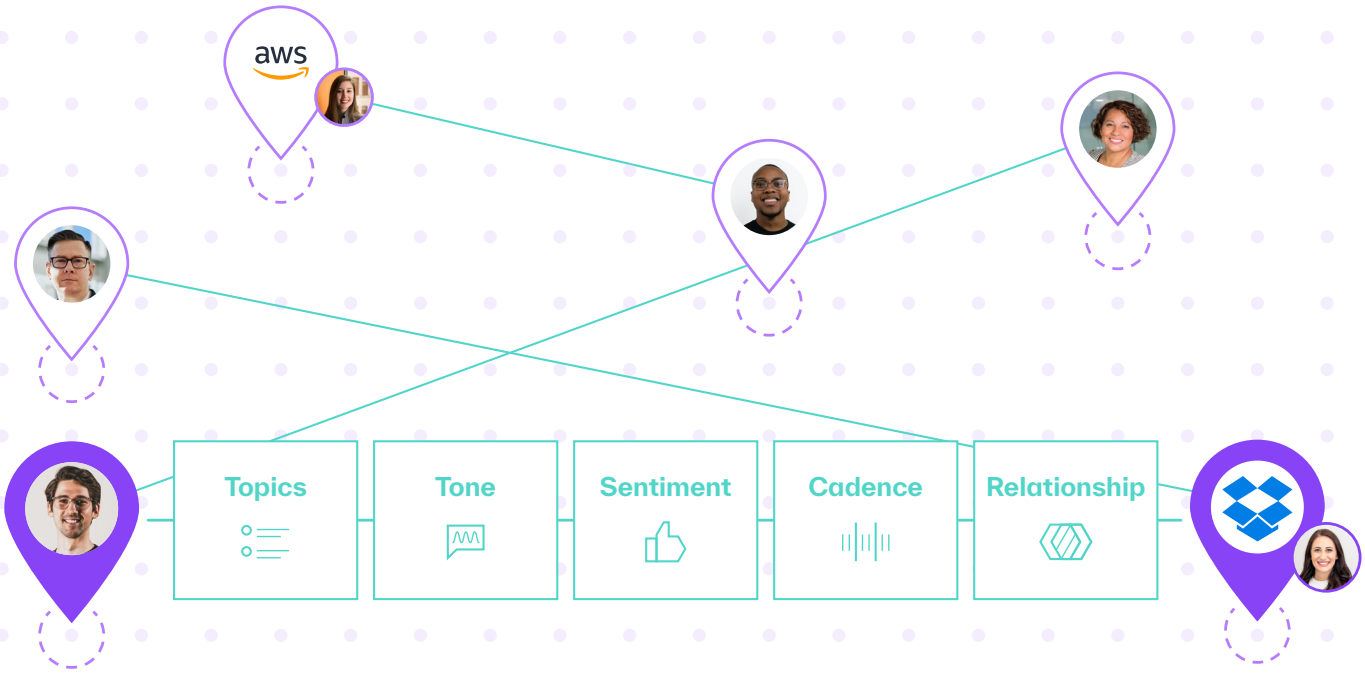

ABX understands the strength of each connection, the structure of the organization, and the frequency of communication among identities, along with the intent, sentiment, and tone of each email communication, making it context aware.

Abnormal Risk Awareness

The risk of each event in the email platform is analyzed by ABX using a variety of techniques, including:

Deep URL Analysis

Abnormal uses behavioral AI to assess the risk of every URL in inbound emails based on its structure, reputation, and message context. URLs contained within attachments are also analyzed.

Extraction Techniques

Abnormal scans and extracts text from within images and other attachments to inform the intent of the message. Attachments are also scanned for malware and the platform detects malicious signals such as macros, executables, javascript, and password protection associated with the files.

Natural Language Algorithms

Natural language understanding (NLU) algorithms identify topic, tone, and sentiment within all communications. Costly business email compromise attacks typically feature urgent requests on financial topics, so identifying these types of communications can assist in the accurate detection of attacks.

Natural language processing (NLP) algorithms are also used to help establish the contextual relationships that explain and score the types of communication, such as formal vs. informal, that are occurring between individuals, departments, and organizations.

Threat Intelligence

In addition, ABX leverages a threat intelligence API and extracts key traditional indicators of compromise including links, domains, and sender attributes. These signals provide additional insight into attacks to help ABX make a final decision on whether an email is malicious.

ABX leverages NLU to understand risks, NLP to explain its insights, and pairs threat intel with the analysis of each event in the email platform to determine the level of risk, making it risk aware.

INSIDE ABNORMAL BEHAVIORAL AI \\ **04**

Composite Analysis

An ensemble of machine learning algorithms evaluates the signals generated by the trio of perspectives from the pillars of Identity Awareness, Context Awareness, and Risk Awareness. By doing so, the algorithms identify specific types of attacks and techniques through multiple algorithms, which results in the delivery of a final email disposition alongside clear, concise, and explainable insights for the security analyst to review.

Most solutions that leverage machine learning technologies result in “black-box” outputs. Some results make sense. Others may not, but security analysts have no mechanism for understanding why and how the algorithms reache a specific conclusion.

In contrast, the Abnormal decision engine explains and summarizes the automated analysis of thousands of signals that were used to detect the attack, providing a full analysis overview with details on why the email was blocked or removed and summed up in an attack score.

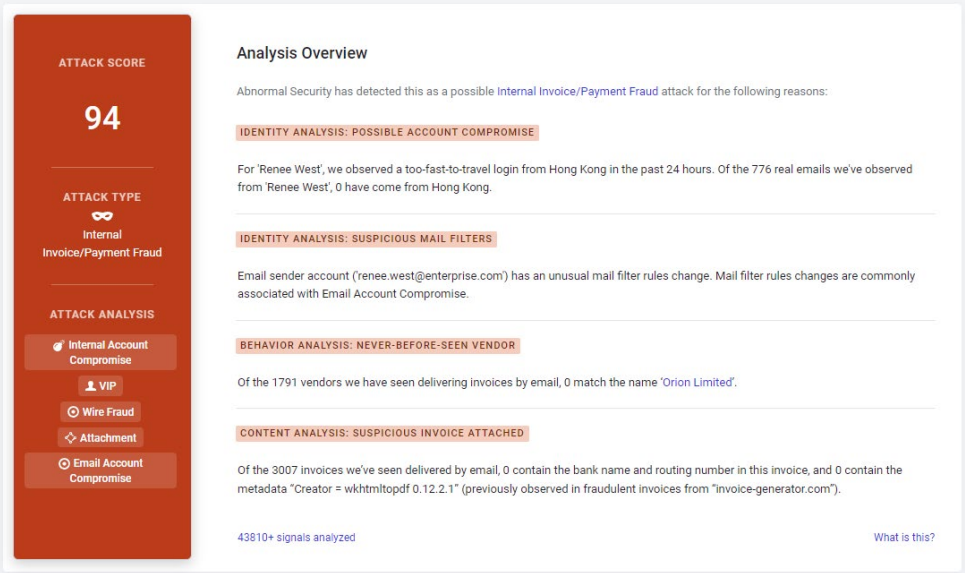

Automated analysis and attack classification provide a clear overview to assist with next steps.

With these insights across identity, context, and risk, SOC analysts can more quickly take action to remediate attacks, address downstream impacts, and educate the organization in real-time about emerging attack types.

The Abnormal Cloud Email Security Platform

Abnormal Behavior Technology powers the Abnormal platform to protect organizations with complete cloud email security. The solution is designed to augment Microsoft 365 and Google Workspace, allowing organizations to reconsider legacy security solutions that cannot stop modern attacks.

The native security capabilities of Microsoft 365 and Google Workspace handle the widespread threats, including broad spam and phishing campaigns, while Abnormal uses its unique behavioral approach to address the sophisticated, targeted inbound attacks and more advanced email platform attacks. Combined, the two platforms can stop the full spectrum of email threats.

The Abnormal Cloud Email Security platform provides six core capabilities:

01. Inbound Email Security

Stops the full range of email attacks, with a unique focus on modern social engineering attacks like business email compromise.

02. Email Account Takeover Protection

Looks beyond email and analyzes hundreds of signals to accurately detect and remediate compromised accounts.

03. AI Security Mailbox

Assists security operations teams with automation and tools to respond quickly to email threats.

04. Email Productivity

Filters time-wasting emails from employee inboxes with an adaptive and policy-free approach.

05. Security Posture Management

Discovers and provides visibility into misconfiguration risks across the entire cloud environment.

06. Misdirected Email Prevention

Reroutes misaddressed outbound emails to quarantine, preventing data loss through AI-powered detection and remediation.

API integration gives Abnormal access to the rich data needed to baseline behaviors and monitor intra-domain traffic—including internal messages that traditional email security cannot see. The API-based architecture also simplifies deployment and maintenance, requiring no MX record or mail flow changes.

Powered by ABX, the Abnormal Cloud Email Security platform protects organizations with cloud-native email security designed to augment Microsoft 365 and Google Workspace.

Conclusion

Abnormal Behavioral AI delivers high-confidence detection of the most advanced socially engineered attacks and vulnerabilities in cloud email. Unlike legacy defenses, it learns from each customer environment, ingesting thousands of identity, context, and risk signals to establish a precise baseline of known-normal behavior. This behavioral intelligence allows Abnormal to detect anomalies that reveal threats such as business email compromise, account takeover, and malicious third-party activity—attacks that evade traditional controls.

By combining identity awareness, context awareness, and risk awareness, Abnormal looks beyond message content to analyze the entire email platform. The result is a security layer uniquely equipped to stop never-before-seen threats, remediate compromised accounts, and provide lasting protection as attackers adopt new AI-powered tactics.

About Abnormal AI

Abnormal AI is the leading AI-native human behavior security platform, leveraging machine learning to stop sophisticated inbound attacks and detect compromised accounts across email and connected applications. The anomaly detection engine leverages identity and context to understand human behavior and analyze the risk of every cloud email event—detecting and stopping sophisticated, socially-engineered attacks that target the human vulnerability.

You can deploy Abnormal in minutes with an API integration for Microsoft 365 or Google Workspace and experience the full value of the platform instantly. Additional protection is available for Slack, Workday, ServiceNow, Zoom, and multiple other cloud applications. Abnormal is currently trusted by more than 3,200 organizations, including over 20% of the Fortune 500, as it continues to redefine how cybersecurity works in the age of AI.

Secure Your Enterprise With Behavioral Intelligence

Request a Demo ›