Abnormal

CISO Guide to Account Takeover

Preventing the Weaponization of Trusted Email Accounts

The Rising Threat of Account Takeover Attacks

60%

chance of a successful account takeover each week for organizations with 50,000+ employees.

Abnormal Security

Compromised accounts in the cloud environment (often referred to as cloud identities) may be the most dangerous threats that organizations face, as they provide cybercriminals with unparalleled access to company data. Once an account—whether on the email platform, SaaS application, or elsewhere in the cloud environment—has been compromised, it can be used to send additional attacks, steal funds and sensitive information, and establish persistence by making changes to key email platform configurations.

Further, account takeovers are hard to detect as attackers are using legitimate login credentials. These are often obtained through successful phishing campaigns, purchased, or scraped from a previous breach. In some instances, attackers don't need credentials at all, hijacking an active session on an application through the use of stolen session tokens.

Once they have access, attackers can move laterally across various platforms, infiltrate critical systems, and send malicious internal email or chat app attacks. Because they appear to be the compromised employee, executive, or vendor, they can establish far more credibility and more easily bypass traditional security measures.

Taking over accounts and compromising cloud identities has proven to be a handy multipurpose tool for cybercriminals. Unfortunately, security teams do not feel well-equipped to stop these attacks with 86% of security leaders and practitioners indicating that current tools are ineffective at stopping account takeover attacks.

How Account Takeovers Begin

Most account takeovers start with a successful login, which requires valid credentials. Phishing, credential stuffing, and brute force password cracking are three ways bad actors can identify the email addresses and passwords they need to hijack email accounts. Increasingly, however, attackers are employing session hijacking to access an active account session without credentials before utilizing social engineering tactics to establish total control.

Phishing for Credentials: Social Engineering at Its Best

Phishing attacks aim to harvest credentials from their targets by impersonating trusted brands, vendors, partners, or executives. By sending an "urgent" message, phishing attacks can trick email recipients into visiting a fake website that logs their credentials as they key them in. No matter who is impersonated in these attacks, the combination of trust and time pressure is a powerful tool for credential theft.

Stuffing Stolen Credentials: Trial and Error at Scale

Sometimes the problem isn't getting access to credentials but figuring out where to use them. Attackers who have a file full of credentials exposed in a data breach can try "stuffing" them into different login pages until they find matches. This may sound tedious, but botnets and AI make credential stuffing fast and scalable.

The volume of available credentials fuels this too. For example, the employee email credentials for 25% of the S&P 500 are among the billions of credentials available on the dark web. And with 81% of people reusing passwords across accounts, that means attackers can break into multiple accounts with the same credentials if they know where to start.

Brute Force and MFA Fatigue: Two Peas in a Password-Cracking Pod

If credential stuffing is like trying stolen keys in every front door on a block, then brute force attacks are like working on one lock relentlessly until it fails. Brute force password cracking attacks use bots and algorithms to generate guess after guess until they hit on the right combination of login ID and password to break into an account. While it's estimated that a highly complex password (12+ letters, numbers, and a special character) could take a computer more than 30,000 years to crack, year after year, the most commonly used passwords are **123456**, **qwerty**, and **password**.

And then there's MFA fatigue. If brute force is persistent lockpicking, MFA fatigue is an attacker ringing the same doorbell over and over again until someone finally answers. MFA fatigue occurs when an attacker logs in with stolen credentials and then, to bypass MFA, sends repeated push notifications to the legitimate user's device. This usually occurs in the early morning or late at night, causing a flustered or groggy user to eventually accept the request to make the prompts stop.

Session Hijacking: No Stolen Credentials Required for Admission

These days, attackers are finding that it's compromise first and get the credentials later. With session cookies for sale on the dark web, malware that can scrape session tokens from browsers, and even complete token forgery, attackers can jump right into an active login session—often for a collaboration app like Slack and masquerade as the legitimate user.

Once a session has been compromised and access has been granted, the next step is often to contact the IT department to request a password reset or a new MFA device be added. Why? Because once a session is terminated, the attacker risks getting booted from the account. But with a bit of clever social engineering, credentials can be reset to ensure continued control long after the legitimate user has ended the session.

How Compromised Accounts Are Used

No matter how cybercriminals access the login credentials, the end result is the same: the compromise of an employee, executive, or trusted vendor's cloud identity—most often focused on a user's cloud email account. And regardless of the method used to snag the credentials, even one successfully compromised account can start a cascade of other internal and external attacks.

Send Lateral BEC and Phishing Attacks

When an attacker with credentials assumes the victim's email identity, they can see all the information in the email account, hijack ongoing threads, and send new email attacks to people on the victim's contact list. When the victim is an executive, the attacker also has the "authority" to direct employees to pay fake invoices, shift the victim's direct deposit to a new bank account, and share insider information for resale, ransom, or corporate espionage.

Access and Manipulate Sensitive Applications and Systems

With the widespread adoption of single sign-on (SSO) solutions over the last five years, organizations have made accessing critical resources more convenient for employees. Although SSO enhances security, a compromised email account can give attackers access to a vast network of critical applications and systems. For instance, a compromised Microsoft 365 account can lead to unauthorized access to Teams, SharePoint, OneDrive, collaboration apps like Slack or Zoom, cloud applications containing sensitive data like Workday and Salesforce, and infrastructure platforms like Amazon Web Services (AWS).

If the victim is a VIP or IT administrator with elevated privileges, attackers can manipulate conditional access policies and reset user passwords to compromise additional accounts.

Create Third-Party Vendor Fraud Attacks

When attackers take over vendor email accounts, they can send fraudulent invoices and requests to update payment account information to any customer of that vendor. Unlike similar bogus requests sent from outside the company's vendor ecosystem, these messages use the same email and invoice formats as legitimate vendor communications. Because they come from a known contact, recipients may not hesitate to make the payment or update account information. Known as vendor email compromise, this tactic is increasingly popular, with fake invoices discovered by Abnormal requesting up to **$36 million**.

Because traditional email security tools don't scan internal, east-west email traffic, they can't detect compromised internal accounts. And because vendor fraud attacks appear to come from legitimate accounts, the recipients are unlikely to question the requests—making these compromised accounts extremely dangerous.

Impact of Account Takeover Attacks

Account takeovers are among the most common tools of the cybercriminal trade, so it is no surprise that 77% of security leaders surveyed in a 2024 analysis indicated that stopping account takeovers is a top priority for their team. In the last year alone, 83% of those same surveyed organizations experienced an account takeover—and half of those attacked organizations experienced an account takeover more than five times in the same year.

This has not only resulted in data and financial loss but also loss of trust in current solutions, as over 60% of respondents were now skeptical of the efficacy of tried-and-true protections such as MFA—and 86% felt all current tools fell short of stopping account takeovers.

Number of Phishing Attacks Reported to FBI IC3

Source: 2022 FBI Internet Crime Report

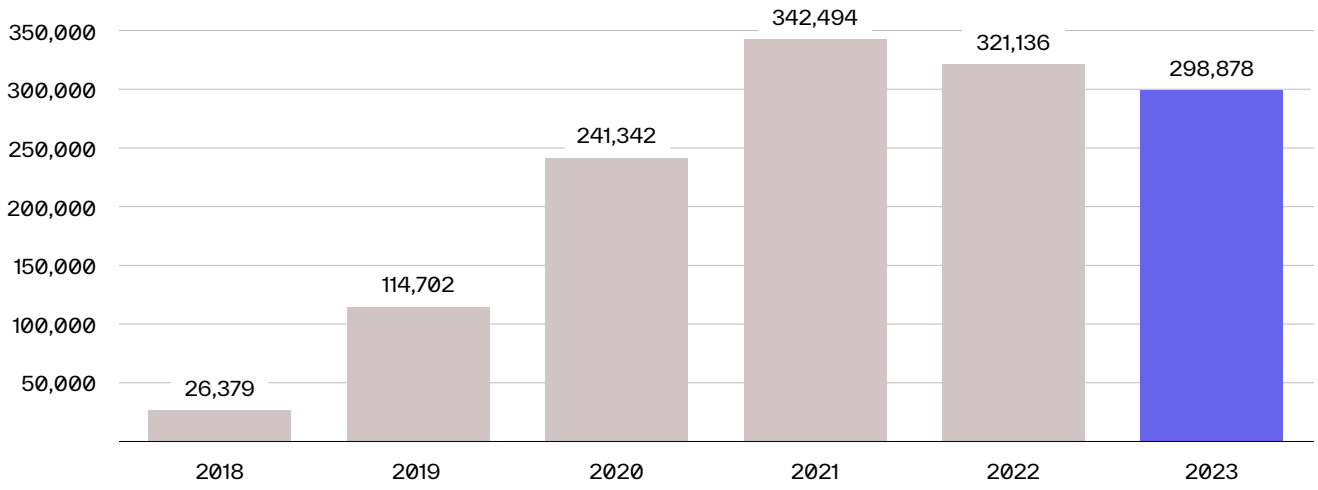

27%

increase in attacks using compromised credentials as the primary vector for cloud intrusion.

IBM Security

66%

of all advanced email threats contain a credential phishing link.

Abnormal Security

775 million

credentials are currently for sale on dark web marketplaces.

Dark Reading

Why Account Takeovers Succeed

Attacks sent from compromised vendor accounts exploit trusted identities and relationships to manipulate recipients' behavior, while compromised internal users are careful to cover their tracks to avoid detection. Identifying compromised accounts requires continuous analysis of behavioral signals that might not be obvious—or even visible.

Current Solutions Meant to Detect Suspicious Sign-In Behavior Often Miss Account Takeovers

While solutions like CASBs and SIEMs are critical pieces of security infrastructure, they are not specifically designed to detect account compromise.

CASBs often rely on policies or known-bad indicators of compromise (IOCs) and may miss attacks that bypass authentication entirely such as session hijacking or from attackers that appear legitimate and do not trigger a preconfigured rule. The SIEM, on the other hand, can ingest signals from across the cloud environment but requires extensive tuning and is prone to noise, necessitating an additional solution that can feed high-fidelity account takeover detections into the SIEM.

Internal-to-Internal Attacks From Compromised Accounts Fool Secure Email Gateways

Secure email gateways (SEGs) can identify known bad senders and screen for indicators of compromise. But because compromised emails come from an ostensibly trustworthy source, and because internal mail sent from a compromised employee typically isn't screened, organizations need more than a SEG to keep attacker messages out of inboxes. Further, SEG solutions often lack additional behavioral signals such as sign-in activity, configuration changes, and privilege escalation, while identity tools lack communication patterns and relationship graphing—all of which are required to confidently detect a potential compromise.

Social Engineering is More Sophisticated Than Ever

Once they have access to an account, bad actors know that the key to further successful attacks is convincingly impersonating the compromised account's owner to minimize suspicion about their requests. Adding the fear of negative consequences can drive immediate action. For example, an urgent email request from the CFO asking to pay a new vendor right away might be unusual. However, if the CFO claims the payment needs to happen before the end of the day to prevent a "contract breach," the recipient is likely to act quickly, without taking time to ask questions.

Security Awareness Training is One Layer Among Many

Security awareness training can help employees avoid opening risky links or attachments by teaching them to spot clues that indicate potential phishing and fraud attempts. However, it's human nature to be less suspicious of messages from senders we know and trust—especially if those messages appear to be urgent requests from more senior employees or crucial vendors. That's why it's critical to supplement security awareness training with technology that can detect messages from compromised accounts.

A Real-World Attack Example

If you look at a real-world example of an attack that bypassed the SEG (as well as MFA and other solutions that may have been in place to detect unusual activity like CASB), you can see why traditional defenses fail. This real-world example, shown via a behavioral timeline, illustrates a multi-pronged attack that bypassed MFA to compromise a VIP account.

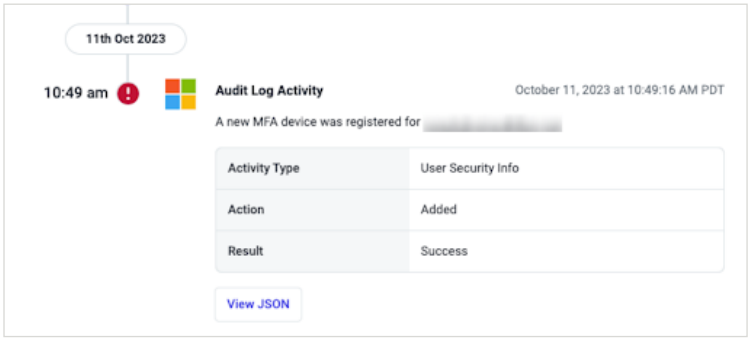

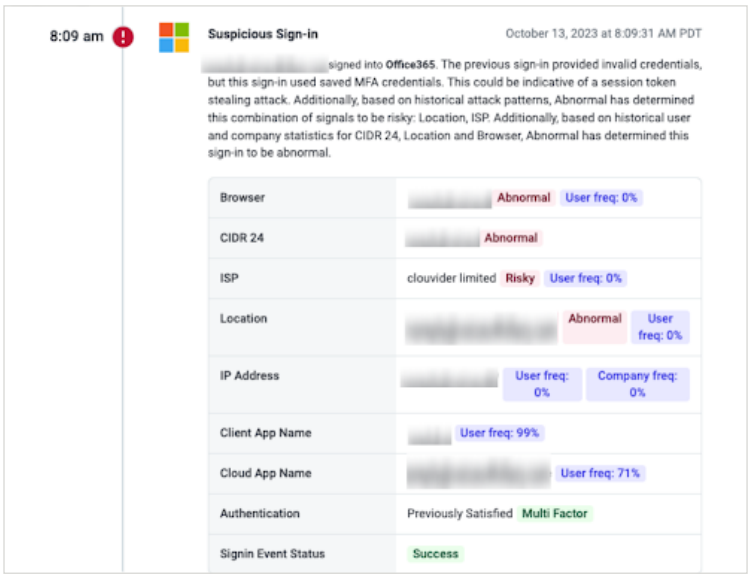

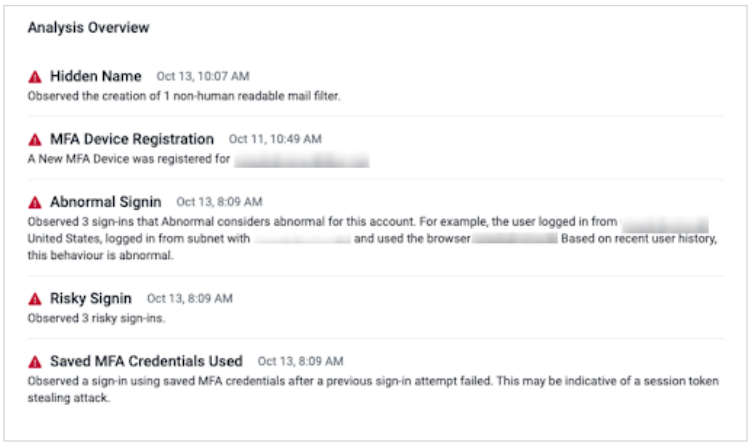

The attack started with a phishing campaign, and while it is unclear whether this was how VIP credentials were initially stolen, the attacker was able to easily infiltrate the account. Likely, as is noted in the analysis, this attacker stole or otherwise acquired an active session—allowing them access to an existing session as a result of the saved MFA credentials.

While this could also be indicative of a legitimate user simply using saved credentials, the analysis further indicates the new sign-in session was from a browser, IP address, location, and ISP that has never been used by the user or the organization. Further, a new, unknown MFA device was registered after this suspicious sign-in, indicating an attacker registering their own device to establish persistence in this account.

Upon gaining this access, the attacker could review all emails and information within the account, move laterally through connected applications, and use the account to send attacks to other employees within the organization, customers, and vendors.

How to Stop Account Takeovers

Protecting organizations from compromised accounts requires security solutions that go beyond just scanning inbound messages for malicious payloads. The next generation of email security is human behavior security and includes:

Multi-Channel Analysis to Benchmark Good Sender Behavior

An API integration with Microsoft 365, Google Workspace, SaaS applications such as Salesforce and Workday, and cloud infrastructure such as AWS and Azure enables the solution to ingest thousands of behavioral signals. From there, the solution should use AI to automatically and continuously analyze communication behavior, login patterns, devices and browsers used, apps accessed, and changes made to user privileges, among thousands of additional factors. The AI engine should quickly learn what normal behavior looks like, create a baseline for each end user, and then analyze anomalous activity to determine whether or not an account has been compromised.

Remediation Options for Compromised Accounts

When user behavior changes, it can be a sign of a compromised account. So, the solution must provide always-on monitoring that uses behavioral AI to look for unexpected changes in user activity, such as changes in content and tone, attempts to bypass multi-factor authentication, and/or shifts in normal login signals. When these events occur and compromise is confirmed, the solution should have the option to rapidly respond by signing users out of active sessions, instantly disabling accounts, and triggering password resets.

Vendor Monitoring to Detect External Compromised Accounts

To detect and prevent compromised external accounts from targeting your organization, the solution should also continuously monitor vendor-customer communication to set behavioral benchmarks and conduct real-time risk assessments. By doing so, it can protect organizations from compromised external accounts that are being used by threat actors to target your organization.

Because these attacks exploit trusted email accounts and relationships, organizations need an email security solution that takes the entire user into account, detecting even small shifts in activity and content. As fraudsters deploy more sophisticated messaging techniques, accurately identifying those minor tells may be the only way to prevent the costly data breaches and financial fraud that can result from a single compromised account.

Conclusion

Cybercriminals aren't likely to give up launching BEC and VEC attacks any time soon, particularly when they can make these attacks more successful with the use of a compromised account. Both internal and third-party account takeovers are relatively easy to execute, and incredibly hard for secure email gateways and identity and access management tools like CASBs to detect. This means that organizations relying solely on legacy email security solutions and current protections like MFA will remain at high risk for costly invoice and billing fraud, data breaches, ransomware, and other high-profile attacks that result from the access that one cloud account can provide.

Detecting compromised accounts and their corresponding attacks calls for a solution that monitors and analyzes thousands of signals through an API to identify indicators of compromise in internal and external emails as well as user behavior across the cloud email platform. Differentiating between good behavior and criminal activity is the most effective way to keep messages from compromised accounts out of end users' inboxes and keep compromised users from executing disastrous breaches.

Abnormal

Abnormal Security 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, Salesforce, ServiceNow, Zoom, Amazon Web Services and multiple other cloud applications.

Ready to Prevent Compromised Accounts?

Request a Demo →

See Your ROI →