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 Research

26%

of companies are targeted by account takeover attempts each week.

Abnormal Security Research

Compromised accounts may be the most dangerous email threat that organizations face, as they provide cybercriminals with unparalleled access to company data. Once an account has been compromised, it can be used to send additional attacks, to steal funds and sensitive information, and to establish persistence by making changes to key email platform configurations.

The takeover of the email accounts is hard to detect because unlike most email attacks, account takeovers happen after criminals receive legitimate login credentials. These are often obtained through the use of phishing messages, brute force attacks, the purchase of stolen credentials, or the attainment of session tokens for session hijacking.

Once they have access, attackers can send malicious email attacks from the actual email accounts of compromised employees, executives, and vendors—establishing far more credibility and more easily bypassing traditional security measures. Further, with full access to email accounts—which often act as the master key to an organization's connected applications and systems—attackers can move freely through critical platforms. These attackers can not only generate difficult-to-detect internal email attacks, but they can also compromise additional communication platforms, change critical security configurations, and access and steal sensitive data across thousands of applications.

Because taking over email accounts has proven to be a handy multipurpose tool for cybercriminals, it's not surprising that 26% of companies receive at least one account takeover attempt each week. And for enterprises with more than 50,000 employees, each week brings a 60% chance that an account takeover will succeed, opening the door to expensive, trust-breaking fraud and security incidents.

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 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 Email 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 email account. And regardless of the method used to snag the credentials, even one successful 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 in the last half decade, organizations have made it more convenient for employees to access critical resources. And while SSO provides a layer of security, a compromised email account gives attackers access to a vast, connected environment of critical applications and systems. A compromise of a Microsoft 365 email account means attackers can likely access Teams, SharePoint, OneDrive or integrated apps like Slack or Zoom from which to survey the organization or exfiltrate data.

Further, when the victim of compromise is a VIP or IT administrator with elevated privileges, this can allow an attacker to directly manipulate everything from conditional access policies to resetting user passwords in an effort to compromise additional accounts.

Create Third-Party Vendor Fraud Attacks

When attackers take over email accounts that belong to vendors, they can then 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 often use the same email and invoice formats as real messages from the vendor.

And because they also come from a known contact, recipients may not think twice about making the payment or account information update. 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 internal compromised 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 one of the most common tools of the cybercriminal trade, with 38% of consumers becoming victims of a successful takeover in 2019 and 2020. This percentage increased in 2021, spiked 131% in 2022, and is likely to continue its unfortunate upward trajectory into the future. In 15% of recent cases, attackers changed the account contact information to lock out the victim entirely and assume control of transaction follow-up and authentication.

When the Verizon Data Breach Investigations Report analyzed 2,249 social engineering account takeover incidents in its 2022 report, it found that 89% of actors had financial motives like invoice fraud and payroll diversion. Among the social engineering takeovers that led to data compromise, 63% exposed credentials and 24% stole personal data.

And when a phishing attack successfully yields credentials, the business email compromise (BEC) and vendor email compromise scams begin. Successful BEC attacks cost organizations close to $2.7 billion in 2022, and Abnormal found that 60% of all organizations receive at least one attack from their supply chain each quarter.

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

Number of Phishing Attacks Reported to FBI IC3

Source: 2022 FBI Internet Crime Report

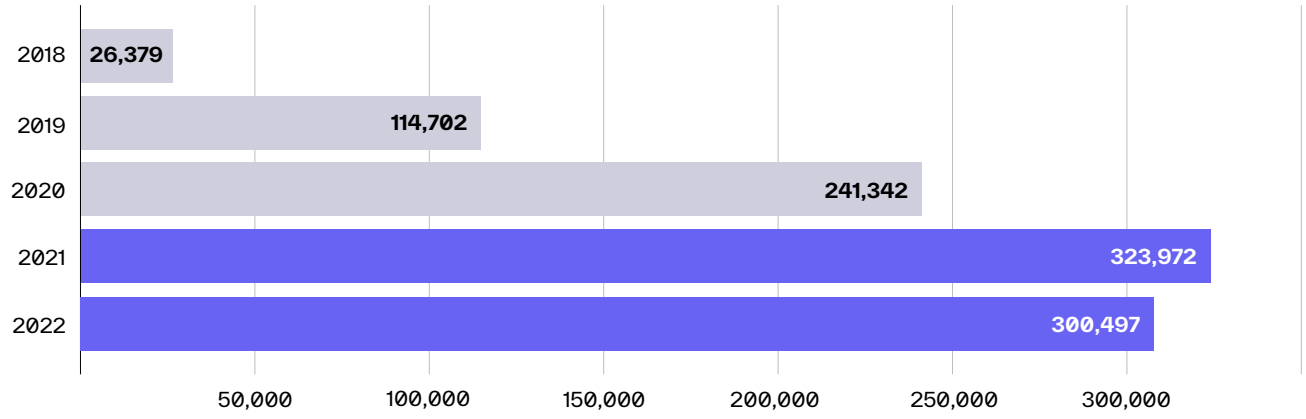

Why Account Takeovers Succeed

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.

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 (and MFA), 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.

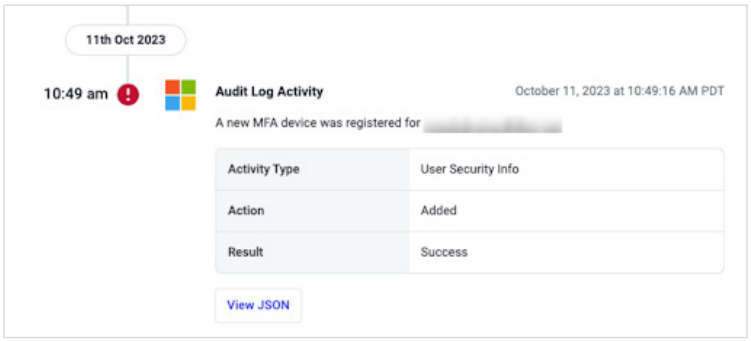

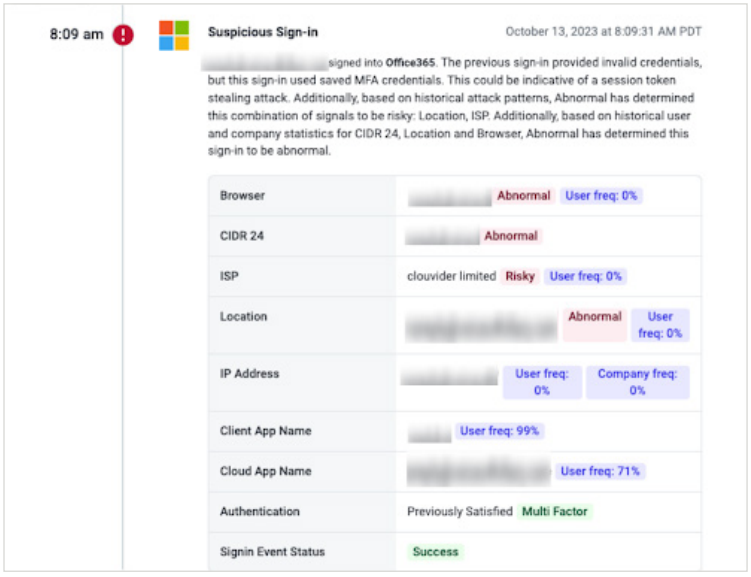

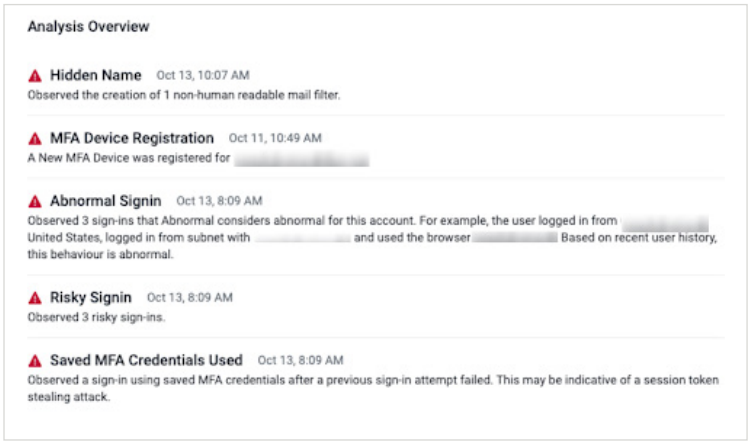

The attack started with a phishing campaign, and while it is unclear whether this was how VIP credentials were initially stolen, the attacker was able to easily infiltrate the account. Likely, as is noted in the analysis, this attacker stole or otherwise acquired an active session—allowing them access 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 then review all emails and information within the account, move laterally throughout connected applications, and use the account to send attacks to other employees within the organization, as well as 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 includes:

Multi-Channel Analysis to Benchmark Good Sender Behavior

An API integration with Microsoft 365 and Google Workspace enables the solution to ingest thousands of behavioral signals. From there, the solution should use AI to analyze communication behavior, login patterns, devices and browsers used, apps accessed, and changes made to user privileges, among thousands of additional signals. 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 external compromised 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 externally-compromised 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 to detect. This means that organizations relying solely on legacy email security solutions will remain at high risk for costly invoice and billing fraud, data breaches, and other high-profile attacks that result from the access that one email 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 provides the leading behavioral AI-based email security platform that leverages machine learning to stop sophisticated inbound email attacks and dangerous email platform attacks that evade traditional solutions. The anomaly detection engine leverages identity and context to analyze the risk of every cloud email event, preventing inbound email attacks, detecting compromised accounts, and remediating emails and messages—all while providing visibility into configuration drifts across your environment. 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, with additional protection available for Slack, Teams, and Zoom.

Ready to Prevent Compromised Accounts?

Request a Demo →

See Your ROI →