Reveal security

White Paper

Safeguard Your Organization Against Account Takeover Attacks

Safeguard Your Organization Against Account Takeover Attacks

Contents

Summary

One of the most effective ways for a threat actor to execute a breach – and escalate it rapidly – is by stealing the credentials of a trusted identity. While most enterprises are heavily invested in identity and access management (IAM) and other sophisticated account controls like privileged access management (PAM), account takeover remains one of the most effective attack vectors. In fact, it's estimated that 86 percent of web application breaches involve the use of stolen credentials.

Cloud and software-as-a-service (SaaS) application adoption is further increasing the identity security challenge, as security teams contend with:

- A tendency to over-privilege users;
- Inconsistent audit log formats and fidelity;
- Lack of consistent policy controls across applications;
- Frequent configuration management by non-security personnel;
- The need to monitor that service providers are fulfilling their shared responsibility for security.

This paper will explore:

- Common techniques used to compromise the credentials of trusted identities.
- How compromised credentials lead to real-world impacts and business consequences.
- Why these attack techniques are difficult for existing security tools to detect and stop.
- How to create a balanced identity security approach that includes effective prevention and detection measures.
- The critical role of identity threat detection and response (ITDR) in extending detection capabilities beyond the point of authentication to ensure that account takeovers cannot blend in undetected with legitimate SaaS and cloud services usage.

86% of web application breaches involve the use of **stolen credentials.**[1]

[1] Source: "2023 Data Breach Investigations Report," Verizon, June 6, 2023.

Common Account Takeover Methods

One of the things that makes account takeovers so challenging to prevent and detect is that they are executed in many different ways.

Common examples include:

 **Credential Stuffing**	Using automated tools to test stolen usernames and passwords from previous data breaches against multiple websites and services.
 **Phishing**	Sending emails or text messages that appear to be legitimate but contain links or attachments that, when clicked, will take the victim to a fake website or download malware onto their computer.
 **Contexting/Vishing**	Using phone-based social engineering, such as posing as technical support representatives, customers, or other trusted individuals, to gain the victim's trust and access credentials.
 **Security Vulnerabilities**	Exploiting unpatched or zero-day software vulnerabilities to deliver malware or otherwise capture user credentials or hijack active user sessions.
 **SIM Swapping**	Tricking or bribing a mobile operator employee to port the phone number of a trusted identity to the threat actor, enabling the interception of multi-factor authentication (MFA) codes.
 **Deepfake Impersonation**	Using generative AI to impersonate the likeness or voice of a trusted identity to defeat biometric verification measures.

While most organizations already take steps to prevent these and other account takeover methods from succeeding, the reality is that none of these approaches are 100 percent bulletproof. And once a threat actor has access to legitimate credentials, further attack escalation becomes extremely difficult to distinguish from legitimate usage.

Most traditional identity security tools are focused on the pre-authentication phase of the user journey and implicitly trust authenticated users. As a result, threat actors using an authenticated account can use the reputation and trust of the hijacked identity to execute even more sophisticated attack techniques.

The only way to protect against the full array of identity threats is to implement a Zero Trust approach. Security teams must assume that breaches of trusted identities are inevitable and implement capabilities to detect suspicious behaviors performed by these identities inside and across applications before the breach escalates.

How Threat Actors Use Account Takeover to Escalate Attacks

Once a threat actor is able to impersonate a trusted identity and access internal systems, they can advance their attack in numerous different ways, including:

- **Reconnaissance**: Exploring the network to understand its structure and identify high-value assets.
- **Lateral Movement:** Accessing other systems, accounts, or network segments, often seeking higher-level privileges or more sensitive data.
- **Privilege Escalation:** Attempting to gain higher-level privileges by exploiting system vulnerabilities, cracking passwords, or finding misconfigurations.
- **Deploying Malware orRansomware:** Using access to deploy malware, ransomware, or other malicious software with the goal of data theft, system damage, or a complete lockdown of the company's systems for ransom.
- **Data Exfiltration:** Accessing sensitive data and extracting it for various purposes, such as selling it on the dark web, using it for identity theft, or leveraging it for competitive advantage.
- **Setting Up Backdoors:** Creating backdoors that will allow them to re-enter the network easily in the future, even if the original compromised account is secured.
- **Launching DDoS Attacks:** Using access to network resources to launch distributed denial-of-service (DDoS) attacks, either as a distraction or as a primary attack vector.
- **Undermining Security Controls:** Using a compromised account to disable security controls, making the network more vulnerable to future attacks and preventing the detection of their activities.
- **Executing Additional Phishing and Social Engineering:** Using a compromised account to conduct additional phishing campaigns or other social engineering attacks against other employees, customers, or partners, leveraging the trust associated with the account to increase the chances of success.
- **Spreading to Third-Party Networks:** Using a compromised account to jump to supply chain partners through interconnected systems.

> "Conventional identity and access management and security preventive controls are **insufficient** to protect identity systems from attack. To enhance cyberattack preparedness, security and risk management leaders **must add ITDR capabilities**.[2]
>
> -Gartner

[2]Source: "Enhance Your Cyberattack Preparedness With Identity Threat Detection and Response," Gartner, October 20, 2022.

The Business Impact of Account Takeover

When threat actors successfully compromise the credentials of a trusted identity using the techniques described above, the business impacts are often severe.

Adverse business outcomes may include:

- Financial loss;
- Reputational damage;
- Operational disruption;
- Legal and regulatory consequences;
- Intellectual property leakage;
- Resource drain;
- Higher cybersecurity insurance premiums.

328 days

The average time it takes an organization to identify and contain breaches resulting from stolen or compromised credentials, 328 days, is longer than the dwell time for any other attack vector[3]

In many cases, these scenarios may also represent a "material breach" in the eyes of key regulators. For example, in September 2023, the U.S. Securities and Exchange Commission (SEC) introduced a new rule for cybersecurity risk management, strategy, governance, and incident disclosure.

One notable requirement of this rule is that organizations that experience a material breach must notify the SEC (in the form of an 8-K filing) within four business days of becoming aware of the event. In this context, the word "material" is carried over from other areas of SEC oversight and applied to cybersecurity. An incident is considered "material" if a reasonable investor would find the information important in deciding whether to buy, sell, or hold a security, or if it would otherwise alter the overall set of data they are considering when making financial decisions. The breadth of this definition underscores the increasing stakes for detecting, assessing, and responding to breaches quickly.

Tougher Regulations and Guidance Emphasize the Critical Role of Identity Threat Detection

Regulators and standards bodies are mobilizing rapidly to ensure that organizations globally are better prepared to defend against account takeovers and other identity-based threats. Increasingly, this guidance is focused on the need for continuous monitoring and detection of identity threats through more sophisticated techniques such as behavioral analytics and anomaly detection.

[3]Source: "Cost of a Data Breach Report, 2023," IBM, July 24, 2023.

EU Digital Operational Resilience Act (DORA)

Article 10.3, which focuses on detection, mandates the use of mechanisms to promptly detect anomalous activities.

NIST Cybersecurity Framework (CSF) 2.0

DETECT section emphasizes continuous monitoring, including behavioral analytics, for the early detection of cybersecurity events.

CISA Zero Trust Maturity Model

Identity pillar advocates for automated analysis of user activity log types, including behavior-based analytics.

MGM Resorts Breach Highlights How 'Vishing' Threats Are Targeting Privileged Users

In September 2023, MGM Resorts experienced a catastrophic security breach that led to a 10-day disruption of critical computer systems, extensive customer data exfiltration, and an overall estimated financial impact of $100 million.

This incident illustrates the increasingly sophisticated techniques that threat actors are using to target trusted identities, including privileged users.

As enterprises become more effective at stopping phishing attempts, threat actors are evolving their techniques in response. Voice phishing, or "vishing," is a primary example. In MGM's case, a sophisticated hacking group known as Scattered Spider researched the company's privileged users using public data sources like social networks and used this information to impersonate them in calls to the MGM IT help desk. Eventually, they were successful in tricking the company into performing a password reset.[4]

The group then used this privileged access to deploy the ALPHV (aka BlackCat) ransomware across MGM's critical systems, wreaking havoc on business operations and customer experience.

Examples like this illustrate why it's critical even for companies with sophisticated IAM approaches in place to implement additional measures to monitor for abuse of trusted identities.

[4]MGM Resorts Hackers Broke In After Tricking IT Service Desk," Bloomberg, September 15, 2023.

Why Existing Security Approaches Frequently Miss Account Takeovers

Most organizations now have a variety of security measures in place to manage trusted identities and access to internal systems. Most provide significant value, but they are largely preventative in nature. They have severe limitations when it comes to detecting that trusted identities have been compromised, driving the need for an additional layer of detection that addresses areas that IAM, PAM, and other preventative controls don't.

For example, IAM platforms are very mature and are highly effective at protecting and governing account access most of the time. But they are not infallible, and there have been numerous examples in recent years of companies – and even IAM vendors themselves – suffering breaches that circumvent this infrastructure. And while IAM products perform some types of risk-based authentication, these checks are generally based on factors such as location and do not consider specific interactions an identity has with an application post-authentication.

Leading IAM Providers Breached By **Account Takeover Techniques**

Okta and Microsoft, two of the most widely used IAM providers, were both breached using account takeover techniques within several months of one another.

In Okta's case, a trusted service account with permission to view and update support cases was compromised when an employee logged into the service account using a Chrome web browser session that was simultaneously logged into their personal Google profile. As part of this process, the privileged service account credentials were saved to the employee's personal Google profile, allowing the threat actors to target the employee's personal account or personal devices to access the credentials.

For a more detailed account of this incident and its impact, read "Closing the Identity Threat Detection Gap: The Okta Support Unit Breach Revisited" on the Reveal Security blog.

In Microsoft's case, the same Russian threat actor group that executed the infamous SolarWinds attack in 2020 used a password spraying attack to compromise an account on one of Microsoft's non-production platforms. The group then used that account's privileges to gain access to the production email accounts of numerous Microsoft trusted identities.

We cover this incident in detail on the Reveal Security blog as well in "Why Microsoft's Latest Breach is an Identity Threat Detection Wake-Up Call" and "Five Lessons from the Microsoft Identity Breach."

While it's not uncommon for security vendors to suffer breaches, these incidents underscore the importance of combining preventative controls like IAM with identity threat detection.

PAM products have similar characteristics. They empower security teams to implement more granular control over individual identities' privileges for specific applications, which is unquestionably a best practice. They also provide detailed reporting about the privilege usage so that policies can be improved regularly. However, they are limited in their ability to detect in-progress abuse of accounts. So even though PAM is a best practice, it will not necessarily limit the damage if a privileged user is targeted with a successful credential compromise.

For these reasons, adding ITDR capabilities is an emerging best practice that acts as a critical layer of defense against compromised credentials. For example, Gartner notes, "Conventional identity and access management and security preventive controls are insufficient to protect identity systems from attack. To enhance cyberattack preparedness, security and risk management leaders must add ITDR capabilities.[5]

How ITDR Reduces Your Exposure to Account Takeover Attacks

ITDR acts as a critical complement to preventative identity security capabilities like IAM and PAM. While IAM and PAM give security teams the tools to manage access and entitlements proactively and precisely, ITDR provides essential protection against unexpected situations when threat actors circumvent these controls.

But it's important to note that not all vendor approaches to ITDR are the same. To be effective, an ITDR solution must:

- Perform continuous threat detection that extends beyond the initial point of authentication.
- Incorporate application awareness into its identity threat detection approach.

While detection of general usage anomalies like unexpected locations, times, usage volume, and download actions is useful, these techniques miss many types of application-based threats. An ITDR solution must build an understanding of how specific application functionality is normally used by specific identities – or group of identities. This is challenging to accomplish, since it's impractical for security teams to understand the intricacies of how every application works – and how it could potentially be misused by a threat actor. But innovations with unsupervised machine learning now make application-aware identity threat detection possible. By baselining normal application user journeys – including cross-application user journeys – at a much more granular level and detecting anomalous behavior accurately, leading-edge ITDR solutions like Reveal Security can detect abuse of trusted identities – post authentication – that traditional security approaches miss.

[5] "Enhance Your Cyberattack Preparedness With Identity Threat Detection and Response", Gartner Inc., October 2022.

Another critical characteristic to look for in an ITDR solution is ubiquitous coverage of all applications and deployment models. Enterprise applications now take many forms, including custom applications, commercial software, open source technologies, and SaaS applications. And these applications frequently span on-premises infrastructure and cloud platforms. An ITDR solution must work across all of these application types and deployment models and provide a unified view of identity-focused threats.

Implemented effectively, ITDR works in harmony with IAM and PAM technologies to reduce the impact of account takeover attacks. If a threat actor successfully circumvents your identity policy controls, the resulting account usage anomalies will be detected within hours with the right ITDR approach in place. This will empower your security team to act quickly to:

- Contain and recover from rapidly escalating attacks.
- Prevent attempts to move laterally and escalate privileges over time.

In addition, ITDR activities provide valuable insights that can be used to continuously improve IAM and PAM policies, keeping them in step with the evolving identity threat landscape.

Start your ITDR Journey with Reveal Security

Reveal Security is an identity threat detection company that baselines the normal user and identity journeys for your SaaS, custom applications, and cloud services and detects anomalies that indicate abuse of a trusted identity. The Reveal Security ITDR Platform uses patented Identity Journey Analytics™ – rather than manually created rules – to detect abnormal behaviors and application usage. This catches threats that traditional monitoring approaches miss, minimizes false positives, and enables fast response when your critical business applications are abused or misused.

Reveal Security ITDR is also the only solution that can detect anomalies within the context of an entire business process that spans multiple applications. This approach delivers a superior level of context and accuracy in detection and enables the business to respond promptly before a material breach can occur.

Contact us today to learn how RevealSecurity can protect your SaaS and cloud services environment from account takeover attacks while reducing the day-to-day burden on your team.

About Reveal Security

Reveal Security quickly and accurately detects identity threats post-authentication in and across SaaS applications and cloud services. The Reveal Security ITDR platform is the only solution in the market based on patented Identity Journey Analytics™ technology that uses unsupervised machine learning to continuously analyze the activity of human and machine identities in applications and detect anomalies. This approach delivers a superior level of accuracy and context, reduces alert volumes by orders of magnitude, and enables the business to respond quickly before a material breach can occur.

For more information, visit www.reveal.security