



CyberMaxx

TALES FROM THE SOC • VOLUME 3

Tales & Trends from the SOC

Real stories from CyberMaxx's Security Operations Center, and what they reveal about where, why, and how attacks originate.

Every incident in this book started the same way every incident always does: quietly.

INTRODUCTION

A login that looked routine. A search result that seemed legitimate. A phone call from someone claiming to be IT. None of these moments looked like the beginning of a breach, and that's exactly the point.

Our SOC doesn't wait for a tidy, high-confidence alert to act. Analysts investigate on suspicion, connect activity across users and clients, and move from detection to containment without a hand-off. That mindset is what turned each of the following eight engagements from a potential large-scale compromise into a contained, well-understood incident.

This installment of Tales from the SOC does something the previous two didn't: after the stories, we step back and look at the bigger picture. Across these eight engagements, attacks trace back to five distinct points of origin. We break down the incidents by attack vector, the techniques used, and what we recommend every organization do to close the gaps attackers keep finding.

Eight stories from the SOC that show where and how attacks really begin

01 **Six Sign-Ins, One Overlooked Login Method**

03 **When a Search Result Leads to a Terminal Window**

05 **A Compromised Website and a Telegram Connection**

07 **A Voicemail Attachment That Wasn't**

02 **A Fake Verification Code Nearly Opened the Door**

04 **Spotted While Looking the Other Way**

06 **A Spoofed Microsoft Email, Sent to a Thousand Inboxes**

08 **The Fake IT Call That Almost Worked**

Then a data-driven look at 56 SOC engagements shows where these techniques cluster — and what to do about it.

It started with a single confirmed account compromise and a wave of outbound email, but a deeper look revealed the threat actor had been quietly working the same technique against multiple users for weeks.

A Familiar Sign-In Technique Turns Malicious

TRT was engaged after confirming a user's account had been compromised through Device Code authentication, a legitimate Microsoft sign-in flow that threat actors increasingly abuse to bypass typical MFA prompts. Once inside, the threat actor created a hidden inbox rule and used the account to send roughly 600 emails, the large majority to external recipients. Four internal users received the message, though no unusual login activity was observed on their accounts during the investigation.

Pulling the Thread Further Back

Looking across the prior 30 days, TRT identified six total users with Device Code driven logins originating from the same IP address, with the earliest login dating back nearly a month before this confirmed compromise. Two of the affected users had already been remediated through separate, earlier incidents, but this investigation revealed the activity was part of one broader, ongoing pattern rather than isolated events.

WHY THIS MATTERS

Device Code phishing succeeds because it looks like a normal sign-in step, no fake login page, no obvious red flag, just a code the user is asked to enter. Without correlating logins across users and time, this kind of slow-burn compromise can persist for weeks. It's a clear example of why identity telemetry, not just email or endpoint alerts, is essential to catching account takeover early.

Six Sign-Ins, One Overlooked Login Method

A single click while browsing the web triggered a textbook ClickFix attack, but fast containment kept it from ever leaving one machine.

A Familiar Trick: The Fake Verification Prompt

TRT was engaged after malicious commands were detected running on a client endpoint. Investigation confirmed the activity matched the ClickFix pattern: a user browsing the web was socially engineered into copying and running a script, believing it to be a routine verification step, which then reached out to a remote server.

Testing the Waters

Once executed, the threat actor ran a basic enumeration command searching for local administrator accounts. A common early step used to gauge what access had just been gained.

Contained Before It Could Spread

Further investigation found no evidence of lateral movement or data exfiltration. The incident stayed confined to the single endpoint where it started. The client isolated the machine, rotated credentials for anyone who had used it, and opted to decommission the device entirely rather than attempting to clean it.

WHY THIS MATTERS

ClickFix attacks rely entirely on a user's willingness to paste and run something they don't fully understand. There are no malicious attachment and no obvious phishing email, just a believable prompt on a compromised or malicious webpage. Fast isolation kept a single bad click from becoming a network-wide incident.

**A Fake
Verification
Code Nearly
Opened the Door**

A routine web search became the opening move in a ClickFix attack, one that endpoint protection stopped just before it could fully execute.

Poisoned Search Results Set the Trap

A user landed on a compromised website through SEO poisoning, a technique where threat actors manipulate search rankings so a malicious site appears among legitimate results. The site displayed a fake “verification code” and instructed the user to paste it into their terminal.

From Copy-Paste to PowerShell

Once pasted, the code launched a PowerShell session that reached out to retrieve a malicious payload. The payload then attempted to execute in a second PowerShell session, but endpoint protection intervened and blocked it before further exploitation could occur.

WHY THIS MATTERS

This incident shows how far upstream an attack can start, not with an email, but with a manipulated search result. It also shows the value of layered defenses: even when a user is successfully socially engineered, strong endpoint protection can still be the difference between a contained event and a full compromise. The host was isolated, the user's password reset, and the machine reimaged out of caution.

**When a Search
Result Leads to a
Terminal
Window**

This compromise wasn't found through an alert within this client's own environment at all. It surfaced while the SOC was investigating a completely unrelated incident for a different organization.

An Unexpected Discovery

While working a separate investigation, analysts broadened their search to make sure this ClickFix-style command did not affect other clients. With this extra effort, analysts noticed the ClickFix-style command executing on an endpoint belonging to another client entirely. TRT was engaged to determine the scope of the compromise and rule out lateral movement or data exfiltration.

A Deeper Attack Chain Than Expected

The investigation confirmed a full ClickFix attack chain, including several stages of persistence. Following the malicious PowerShell activity, multiple programs were dropped onto the host, adding autorun registry keys to survive a reboot. The majority of command-and-control traffic was blocked outright because the host had already been placed under network isolation.

Fast Containment Limited the Damage

No signs of lateral movement or data exfiltration were found, and the incident remained confined to a single host. Isolating the endpoint quickly after the malicious command executed significantly reduced the threat actor's dwell time. The team recommended a full reimage, a credential reset, and revoking active sessions.

WHY THIS MATTERS

Cross-client visibility matters. This threat was only caught because analysts were paying attention beyond the scope of the case directly in front of them. It's a reminder that patterns worth investigating don't always announce themselves through the client or alert you happen to be working at the time.

**Spotted While
Looking the
Other Way**

An alert about interactive PowerShell activity led analysts down a multi-stage attack chain that ended with traffic to an app most people associate with messaging, not malware.

An Alert Worth Chasing

An automated alert flagged interactive PowerShell activity alongside signs of a potential penetration testing framework. Investigation traced the activity back to a compromised WordPress site the user had visited less than a minute before the malicious PowerShell execution began, a connection that didn't surface when analysts navigated to the same site afterward in a sandboxed environment.

A Multi-Stage Payload

The attack unfolded in several stages: an installer package pulled from a remote server, additional tooling delivered through a compressed archive, and yet another executable retrieved from a separate remote source. Analysts also observed DNS requests to suspicious domains and to Telegram's API, a platform frequently abused by threat actors for command-and-control and data exfiltration.

Attackers layer legitimate-looking infrastructure to blend in with normal traffic.

WHY THIS MATTERS

Not every compromised website looks malicious after the fact, sandboxing the same site later showed nothing unusual, which is exactly what makes this technique effective. The multi-stage delivery chain also illustrates how attackers layer legitimate-looking infrastructure, including everyday platforms like Telegram, to blend in with normal traffic. The SOC isolated the host and advised the client to block the identified indicators, rotate credentials, and engage digital forensics for a deeper look.

A Compromised Website and a Telegram Connection

One convincing email impersonating Microsoft was the opening move in a Device Code phishing attack that reached roughly a thousand recipients before the compromised account was contained.

Tracing the Spoofed Email Back

TRT determined the compromise stemmed from Device Code phishing, a technique that abuses a legitimate Microsoft authentication flow to trick users into unknowingly granting a threat actor access. Investigators traced a spoofed email impersonating Microsoft that had landed in the user's inbox shortly before the first malicious login, confirming it as the point of initial access. Message trace data showed the same spoofed email had gone out to approximately 1,000 users.

Establishing a Foothold

Once inside, the threat actor registered new devices under the compromised account, a common tactic for maintaining persistent access even after a password reset. The compromised account was also used to send emails to both internal and external recipients. A review confirmed no other accounts experienced successful Device Code logins.

Cleanup and Containment

The client had already reset the account's password and revoked active sessions by the time TRT completed its review. The team recommended the client also redo MFA registration given the newly added devices, recall the emails sent from the compromised account, and block the sender behind the spoofed message.

WHY THIS MATTERS

Device Code phishing is effective precisely because it doesn't look like traditional phishing. There's no fake login page to inspect, just a legitimate-looking prompt to enter a code. When a single email can reach a thousand inboxes, the margin for error narrows fast, which is why rigorous end-user awareness training and tight controls around device code authentication matter as much as any technical control.

A Spoofed Microsoft Email, Sent to a Thousand Inboxes

An external email with an audio-sounding subject line marked the start of a business email compromise that moved from inbox to inbox quickly.

A Convincing Subject Line

The attack began with an external email carrying a subject line designed to look like a routine voicemail notification. Within roughly a day, the SOC observed the first malicious login to the targeted account.

A Rule Built to Hide the Evidence

Shortly after the malicious login, a rule was created inside the mailbox that automatically deleted incoming mail, a technique threat actors commonly use to keep a compromised account under their control without the real user noticing. Email activity showed the compromised account was then used to send out multiple messages. At the time of review, no additional users had interacted with those outgoing emails.

Rapid Response Over a Vulnerable Window

The SOC rotated credentials and revoked sessions for the impacted user, recommended the client recall the emails sent from the account, and advised blocking the sender behind the original phishing email. Given the timing, the client was also placed under increased monitoring vigilance heading into the weekend, when detection and response resources across many organizations are often thinner.

***One convincing email and a mailbox rule.
Defenders had almost no time to react.***

WHY THIS MATTERS

Business email compromise doesn't need malware to succeed, just one convincing email and a mailbox rule. The short window between the phishing email and the malicious inbox rule shows how little time defenders have once initial access is gained. This is a direct example of why advanced email filtering and monitoring, paired with detection tuned to flag mailbox rule creation, are critical controls.

A Voicemail Attachment That Wasn't

A threat actor posing as internal IT reached two employees with the same pretext, and one of them handed over remote access to their machine.

An IT Impersonation Attempt

Two users were targeted by a threat actor impersonating the organization's IT department. One user was convinced and granted remote access to their machine through a legitimate remote support tool.

From Remote Access to Data Staging

Once inside, the threat actor ran a command to enumerate the domain, interacted with the Volume Shadow Copy Service to review snapshots and groups before disabling it, and downloaded a file transfer tool commonly associated with data exfiltration. Investigators identified a handful of connections to an external cloud-hosted address involving a modest amount of outbound data, small enough to raise questions on its own, but consistent enough with the use of a file transfer tool to treat as a likely exfiltration attempt.

Weighing the Next Steps

The client discussed the incident directly with CyberMaxx's digital forensics team but opted to hold off on a full forensic engagement for the moment, preferring to speak with the affected employee first to understand exactly what had been accessed.

Not a link. Not an attachment. Just a convincing phone call.

WHY THIS MATTERS

This incident is a clear example of help desk impersonation, a social engineering tactic that doesn't rely on a malicious link or attachment at all, just a convincing phone call and a legitimate remote access tool. The SOC isolated the endpoint and recommended locking down collaboration tools, rotating credentials, blocking the external destination, disabling the remote access tool where it isn't needed, and reinforcing social engineering awareness training, the same theme that runs through nearly every story in this book.

The Fake IT Call That Almost Worked

Where These Attacks Are Really Coming From: Three Attack Vectors and Five Techniques Behind Eight Incidents

Looking across these eight engagements as a set, the initial access techniques cluster into three distinct attack vectors. Understanding where these attacks actually originate, rather than treating each incident as a one-off, is what makes it possible to prioritize defenses.

ATTACK VECTOR: ENDPOINT

Malicious or Compromised Websites (ClickFix)

The most common attack vector across these stories was the endpoint. In each case, a user was socially engineered while simply browsing the web, tricked into copying and pasting attacker-supplied commands into a terminal or run dialog under the guise of a "verification" or "fix" step. No attachment, no email required.

3 / 8

ATTACK VECTOR: ENDPOINT

SEO Poisoning

The attack came through the endpoint, same as above, but the entry point is a manipulated search result rather than a link a user clicked from elsewhere. Threat actors gamed search rankings, then used the same ClickFix-style technique once the user landed there.

1 / 8

ATTACK VECTOR: EMAIL

Device Code Phishing

Both incidents abused Microsoft's device code authentication flow, a legitimate sign-in mechanism intended for devices without a browser, to gain account access without needing to defeat MFA directly.

2 / 8

ATTACK VECTOR: EMAIL

Business Email Compromise via Phishing

A single well-crafted phishing email led directly to account takeover, a malicious inbox rule, and outbound phishing sent from the compromised account, all within about a day.

1 / 8

ATTACK VECTOR: IDENTITY

Help Desk / IT Impersonation (Social Engineering via Remote Access)

A threat actor impersonated internal IT over the phone and convinced an employee to grant remote access through a legitimate support tool, then used that access to enumerate the environment and stage data for likely exfiltration.

1 / 8

THE TAKEAWAY

Seventy-five percent of the incidents in this book began with some form of social engineering: a fake verification prompt, a poisoned search result, a spoofed email, or a phone call impersonating IT. Identity abuse (Device Code phishing and BEC) accounted for another meaningful share. In nearly every case, the technology worked as designed; a person was convinced to do something they shouldn't have.

A Deeper Analysis Across the Five Attack Vectors: Going Beyond These 8 Stories

Inclusive of the incidents shared in this book, we analyzed a subset of 56 incidents from our TRT and SOC engagements over the past three months to better understand where modern attacks are actually originating. The findings reveal a clear pattern: attackers continue to favor identity- and email-centric techniques, often combining multiple attack vectors throughout a single campaign.

CATEGORY	INCIDENTS	SHARE
Identity	35	62.5%
Email	29	51.8%
Endpoint	23	41.1%
Network	5	8.9%
Cloud	3	5.4%

Each incident was categorized by where the attack originated. Because many incidents involved multiple vectors, totals exceed 56.

Identity: Nearly Two-Thirds of All Incidents

The prevalence of identity-based attacks was driven primarily by credential theft and account takeover activity. Abuse of Microsoft’s device code authentication flow was particularly common, appearing in roughly one in five incidents across the analysis period.

Endpoint: About Four in Ten Incidents

Endpoint compromise was dominated by ClickFix-style social engineering and SEO poisoning. Together, ClickFix and SEO-poisoning-related activity appeared in roughly one-quarter of all incidents analyzed.

Email: Just Over Half of All Incidents

Email remained one of the most effective pathways into an organization. Common patterns included phishing campaigns designed to steal credentials, business email compromise following account takeover, and attacker-created inbox rules used to conceal activity.

Network and Cloud: Comparatively Rare

Network and cloud vectors each appeared in fewer than 10% of incidents. When they did occur, they more often represented later stages: command-and-control, lateral movement, or exfiltration.

Attackers Rarely Rely on a Single Technique

Most notably, attackers rarely relied on a single technique. Two-thirds of all incidents involved multiple attack vectors working together, with email and identity compromise appearing together in nearly half of all cases. This reinforces an important reality for defenders: attacks increasingly move across identity, email, and endpoint environments, making visibility across all three critical for early detection and response.

VECTOR COMBINATION	INCIDENTS
Email + Identity	25
Identity + Endpoint	3
Endpoint + Network	3
Cloud + Identity	2
Email + Endpoint	2
Identity + Endpoint + Network	2

Of the 56 incidents analyzed, 37, exactly two-thirds, involved two or more attack vectors working together.

1. Email plus Identity is by far the dominant pattern, accounting for 25 of the 37 cross-vector incidents. The typical shape: a phishing email or device code phishing attempt compromises a user's credentials, the attacker then creates a hidden inbox rule to cover their tracks, and in many cases uses the compromised account to send further phishing emails, extending the attack to new victims. Email is rarely the whole story. It's usually the opening move in an identity compromise.
2. The remaining cross-vector combinations show attacks moving from identity or email into endpoint compromise (helpdesk impersonation leading to remote access and malware execution) or from endpoint into network activity (malware establishing command-and-control after initial execution).
3. The three-way Identity + Endpoint + Network cases represent the most developed attacks in the dataset: social engineering for initial access, followed by endpoint execution, followed by an attempted or successful network-level command-and-control connection.

This reinforces a pattern already visible in this SOC's caseload: identity-driven compromise, particularly device code phishing, and endpoint compromise via ClickFix/SEO poisoning are the two dominant starting points, and email is the connective tissue that shows up alongside identity compromise more often than it appears alone. Defenses built around any single vector will consistently miss the majority of these engagements.

What the SOC Recommends

Based on the patterns above and the response actions taken across these engagements, we recommend organizations prioritize the following:

Invest in rigorous, ongoing end-user training on social engineering and phishing.

Nearly every incident in this book started with a person, not a piece of malware, being convinced to do something they shouldn't have. Training has to go beyond generic phishing email awareness: fake "verification code" prompts (ClickFix), unsolicited calls from "IT" asking for remote access, and out-of-context device code sign-in requests. Awareness training should be recurring and updated as tactics evolve, not a once-a-year checkbox.

Implement advanced email filtering and monitoring.

Both the business email compromise and Device Code phishing incidents in this book relied on email reaching a user's inbox uncontested. Advanced filtering should extend beyond spam and known-bad senders to include monitoring for mailbox rule creation, especially rules that delete, hide, or forward mail, since that's frequently the first sign an account has already been compromised.

Ensure visibility across endpoint, identity, and email simultaneously.

Several of these incidents were caught, or fully scoped, only because analysts could correlate signals across more than one domain: a login pattern tied to identity, a PowerShell execution tied to the endpoint, a spoofed message tied to email. Ask whether your current tools and provider give you all three, together, not in separate silos.

Disable device code authentication flow unless there is a specific, documented business need for it.

Two of the eight incidents in this book, and a meaningful share of TRT engagements more broadly, trace back to abuse of the device code sign-in flow. If your organization doesn't have that specific use case, disabling device code flow removes an entire attack vector with minimal operational impact.

ADDITIONAL PRACTICES WORTH REINFORCING

- 01 Isolate and investigate quickly.** In multiple stories in this book, the difference between a single-endpoint incident and a much larger one was how fast the affected host or account was isolated after detection.
 - 02 Treat "small" or unresolved indicators as worth chasing.** Several of these incidents began with something an analyst could easily have dismissed: a single recurring alert, activity discovered while working an unrelated case, or a modest amount of outbound network traffic.
 - 03 Restrict and monitor remote access tools.** Legitimate remote-support utilities were used against the organization in more than one engagement in this book. Where these tools aren't needed for business purposes, disable them; where they are needed, monitor their use closely.
 - 04 Rotate credentials and revoke sessions immediately** following any suspected account compromise, and don't assume a password reset alone is sufficient. Newly registered devices and unrevoked sessions were both used to maintain persistence after initial containment in these stories.
-

CLOSING

The threats in this book didn't succeed because of a novel exploit. They succeeded because a person clicked, pasted, or trusted something they shouldn't have, and attackers pressed that advantage across whichever surface gave way first: identity, email, endpoint, network, or cloud.

That pattern holds across every story here. Device code phishing, business email compromise, and endpoint compromise are the most common starting points, and two-thirds of these attacks crossed two or more vectors at once. Add AI accelerating how fast attackers move, and one conclusion follows: no single point solution can keep up.



That's why CyberMaxx provides a layered MDR solution.

A human-led, AI-powered SOC and Threat Response Team that monitors, detects, and fully remediates threats across all five attack vectors. It's the difference between a contained incident and a headline.

Ready to take the first step toward a stronger security posture?

1-800-897-CYBER (2923) • CyberMaxx.com