CSAI Foundation | Cloud Security Alliance

ACR Stealer's ClickFix Campaigns Drain M365 Tokens

2026-07-17

AI-assisted Rapid Research

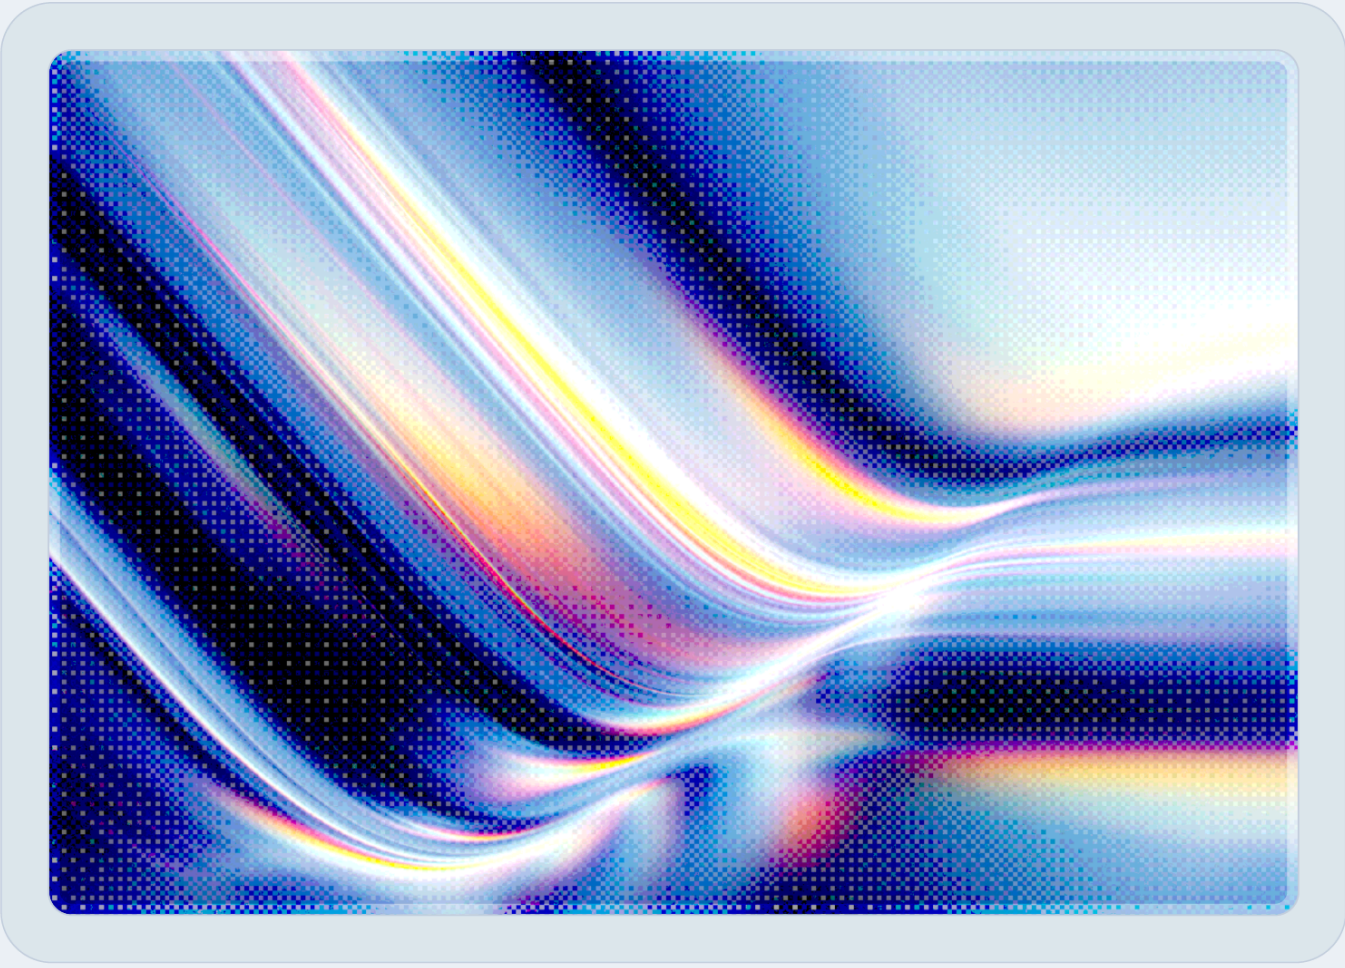

CSAI

CSA cloud security alliance®

Key Takeaways

- Microsoft Defender Experts tracked a sustained rise in ACR Stealer activity from late April through mid-June 2026, delivered through two distinct ClickFix intrusion chains that both begin with a victim pasting an attacker-supplied command into the Windows Run dialog [1][2].
- One chain loads a DLL from a WebDAV share via `rundll32.exe`; the other uses `mshta.exe` to fetch an HTA that decodes PowerShell and retrieves a payload steganographically hidden inside a JPEG, executing entirely in memory [1][2].
- Both chains harvest Chrome and Edge credential and session-token databases via DPAPI, then separately enumerate and exfiltrate PDFs, Microsoft 365 documents, and OneDrive/SharePoint files rather than relying on a single generic "steal everything" routine [1][2].
- A notable subset of infrastructure uses EtherHiding, retrieving C2 addresses or follow-on payloads from public blockchain smart contracts, which lets operators rotate infrastructure without redeploying malware or triggering domain-reputation defenses [1].
- Because compromise requires no software vulnerability, exploitation depends entirely on the user executing a copy-pasted command, meaning patching has no bearing on this attack vector; the controls most likely to break this chain are application control, browser/session hardening, and user recognition of ClickFix-style prompts [1][2].

Background

This assessment treats ClickFix as one of the more consistently effective delivery techniques in the commodity malware ecosystem, because it circumvents common technical controls rather than exploiting a vulnerability in them. Rather than exploiting a vulnerability in a browser, document viewer, or operating system, it exploits the user's willingness to follow on-screen "verification" instructions: a fake CAPTCHA, a spoofed software installer, or a bogus troubleshooting step that asks the victim to open the Windows Run dialog (or a macOS Terminal) and paste a command they do not understand. Because the victim executes the payload themselves, through interpreters and system binaries already present and trusted on the endpoint, ClickFix bypasses code-signing checks, browser sandboxing, and much of the telemetry that endpoint tools use to flag externally delivered executables. CSA's AI Safety Initiative has tracked this technique across multiple prior campaigns, including a macOS-focused infostealer, MacSync, that used the identical lure structure months before the ACR Stealer activity described here [3].

The malware at the center of this campaign, ACR Stealer (also documented under the earlier name AcridRain), has a lineage that illustrates how commodity malware families persist even after their original operators exit the market. The strain was marketed on Russian-speaking cybercrime forums by an actor using the handle "SheldIO," who announced a halt to sales in July 2024 [1]. Threat intelligence vendors diverge on what happened next, as reported by The Hacker News: eSentire has reported that the source code was sold onward, while Proofpoint's account holds that the seller's Telegram channel announced a shutdown rather than a change in ownership, with a related panel, associated with a follow-on strain called Amatera Stealer, surfacing roughly five months later [1]. The Hacker News further cites Proofpoint's mid-2025 finding that ACR Stealer had been substantially updated and rebranded as Amatera, sold on a subscription basis priced from $199 per month to $1,499 per year [1]. Red Canary, per the same reporting, treats ACR and Amatera as a single family lineage and further assesses ACR itself as an update to an earlier strain, GrMsk Stealer [1]. Microsoft's own advisory is notably more conservative on attribution: it does not name a specific threat actor or group behind the activity it observed, instead tying the intrusions to the ACR Stealer family through behavioral analysis and shared post-exploitation tradecraft rather than definitive actor identification [2]. That caution is a useful reminder that malware-family naming in commodity crimeware is often an approximation of a fluid, resold, and rebranded criminal supply chain rather than a stable indicator of a single operator.

The scale of the current wave stands out: Red Canary's telemetry for May 2026 placed ClearFake, the web-injection cluster that feeds victims into this delivery pipeline, at the top of its monthly threat list for the first time, with ACR Stealer itself entering the top ten in a tie for sixth position [1]. The Hacker News also cites a SANS Internet Storm Center handler who documented a Windows infection originating from a Claude-impersonating page reached through malicious Google advertisements on May 26, 2026, corroborating the malvertising vector Microsoft and other researchers describe [1]. Taken together, these signals point to a campaign that has moved well past isolated incident status and into the range of a top-tier commodity threat competing directly with established families like Lumma Stealer for the attention of defenders.

Security Analysis

The Lure: AI Brand Impersonation as the Social-Engineering Hook

The primary infection vector is malvertising and search-engine-optimized results that steer victims toward pages impersonating Anthropic's Claude, frequently hosted behind legitimate-looking `sites.google.com` URLs to evade basic domain-reputation filtering [1]. These pages detect the visitor's operating system and serve platform-specific instructions, Windows or macOS, tailoring the ClickFix prompt to whichever command interpreter is available on the victim's machine [1]. A secondary set of lures uses fake CAPTCHA verification pages and spoofed Claude Code pages hosted on legitimate developer platforms such as GitLab [1]. The choice of Claude as the impersonated brand is consistent with the pattern CSA documented in the MacSync campaigns, where attackers rotated through OpenAI Atlas, ChatGPT, and Claude Code impersonation to capitalize on the surge in developer and knowledge-worker interest in AI coding assistants [3]. The technique's durability across at

least two apparently unrelated malware families over an eight-month span suggests that AI-tool impersonation has become a default lure template for infostealer operators generally, not a one-off tactic tied to a single group, though Microsoft's own reluctance to attribute ACR Stealer to a specific actor means that shared lineage cannot be fully ruled out.

Two Intrusion Chains, One Objective

Both chains observed by Microsoft begin identically, with the victim pasting a ClickFix command into the Run dialog, but they diverge sharply in execution mechanics and forensic footprint. The table below summarizes the distinguishing characteristics of each.

Characteristic	WebDAV/DLL Chain	MSHTA/Steganographic Chain
Loader trigger	`cmd.exe` invokes `rundll32.exe` to load a DLL from a remote WebDAV share over HTTPS	`mshta.exe` fetches remote HTA content containing an embedded VBScript/COM loader
Disk footprint	Writes a ZIP payload to `%LocalAppData%\Temp` under innocuous names (e.g., "LogiOptionsPlus")	Executes reflectively in memory; minimal to no disk artifacts
Payload delivery	Direct download from GUID-named WebDAV directories with obfuscated filenames	Payload steganographically embedded in a JPEG retrieved from an image-hosting service
Execution engine	Bundled `pythonw.exe` runs an obfuscated Python loader that performs shellcode injection via VirtualAlloc and the Windows Fiber API	PowerShell disables certificate validation, generates a victim ID, and resolves LoadLibrary/GetProcAddress/CreateThread at runtime for reflective execution
Stealthiest variant	Wraps execution in `conhost.exe --headless` with delayed environment-variable expansion to hide strings	N/A (chain is fileless by design)
Representative infrastructure	`sphere-api.dialectosphere.in[.]net` (WebDAV host)	`creativecommunityinfo[.]art` (payload host), `enhanceblabber[.]cc` (C2)

Sources: Microsoft Threat Intelligence [2]; The Hacker News [1].

Despite these mechanical differences, the two chains converge on the same downstream objective: obtaining browser-stored credentials and session tokens, then pivoting into enterprise document stores. Both invoke the Windows Data Protection API to decrypt the "Login Data" and "Web Data" databases that Chrome and Edge use to store saved passwords and, critically, active session tokens; stealing a live session token is functionally equivalent to stealing a valid, already-authenticated login, which is why Microsoft's guidance emphasizes token revocation over password rotation alone [2]. From there, the malware enumerates PDFs on the Desktop and in Downloads, along with Microsoft 365 documents and any locally synced OneDrive or SharePoint content, giving the operators a combined haul of authentication material and the business documents that material would otherwise be used to protect.

Persistence and evasion techniques layered on top of both chains read, in this assessment, as consistent with a mature, professionally maintained crimeware kit rather than an opportunistic one-off tool. Observed behaviors include hidden scheduled tasks disguised as routine software updates, timestomping (copying file timestamps from a trusted binary such as `notepad.exe` onto the malware's own files to blend into normal filesystem activity), and deliberate clearing of PowerShell history to remove command-line evidence after execution [1][2]. A subset of the infrastructure additionally uses EtherHiding, a technique in which the malware queries public blockchain RPC endpoints and Web3 node infrastructure to retrieve follow-on payloads or updated command-and-control addresses from smart contracts. Because blockchain state is public, decentralized, and not subject to conventional domain takedown, this gives operators a resilient channel for rotating infrastructure that is materially harder for defenders to disrupt than a traditional DNS-based C2 model.

What Is Actually at Risk in Microsoft 365 Environments

It is worth being precise about the enterprise exposure here: ACR Stealer does not exploit Microsoft 365 itself, and no vulnerability in Exchange Online, SharePoint, or the M365 client applications is involved. The risk is entirely downstream of endpoint compromise. Once a session token or saved credential is exfiltrated from a browser profile, an attacker can potentially reuse it to access whatever cloud services that session was authenticated to, up to and including a user's mailbox, SharePoint sites, and OneDrive, without ever needing the user's password or triggering a fresh multi-factor authentication challenge, since the token already represents a completed authentication event. Locally synced OneDrive and SharePoint folders compound this exposure by placing a copy of cloud-hosted enterprise content directly on a device already known to be compromised, effectively converting an endpoint infection into a document-level breach of cloud collaboration data without any additional attacker effort against the cloud service itself.

Recommendations

Immediate Actions

Organizations that identify indicators consistent with this campaign, or that operate in sectors with elevated exposure to malvertising-driven infostealer campaigns, should prioritize hunting for `rundll32.exe` processes launched with no command-line parameters that are simultaneously making outbound network connections, a pattern Microsoft specifically calls out as a high-value detection opportunity for this activity [2]. Any account with evidence of browser credential-store access from an unexpected process, or with abnormal compression or archival activity in temp directories, warrants immediate session token revocation rather than a password reset alone, since a stolen live token remains usable until explicitly invalidated regardless of password changes [2]. Security teams should also review Entra ID and Microsoft 365 sign-in logs for the affected accounts for anomalous access to mail, SharePoint, or OneDrive originating from unfamiliar IP ranges or device identifiers in the period following suspected compromise.

Short-Term Mitigations

Because both intrusion chains depend on living-off-the-land binaries that are legitimate parts of Windows, technical controls should focus on constraining what those binaries are permitted to do rather than attempting to blocklist the malware itself. Blocking `mshta.exe` outright through AppLocker or Windows Defender Application Control, and restricting PowerShell, Python, and `rundll32.exe` from launching content retrieved from the internet, directly disrupts both chains at their execution stage [2]. Removing or restricting the Windows Run dialog through Group Policy for user populations that do not require it eliminates the initial trigger point entirely for a meaningful share of the workforce. Organizations should also apply web filtering that blocks newly registered or low-reputation domains, since the malvertising and SEO-poisoning vectors described here depend on victims reaching attacker-controlled infrastructure through search and ad platforms rather than direct delivery.

Strategic Considerations

The recurrence of AI-brand impersonation as a lure, first in the MacSync campaigns and now in ACR Stealer's Claude-themed pages, indicates that security awareness programs need to specifically address AI-tool installation and "verification" prompts as a distinct social-engineering category, rather than folding them into generic phishing training that employees may have already learned to tune out. More structurally, the consistent pattern of session-token theft as the actual objective, across unrelated malware families and delivery mechanisms, argues for reducing organizational dependence on long-lived, locally cached browser sessions as the primary access control for sensitive cloud services. Conditional access policies that evaluate device compliance and sign-in risk continuously, rather than only at initial authentication, reduce the practical value of a stolen token even when theft occurs, and should be treated as a baseline expectation for any account with access to Microsoft 365 mail, SharePoint, or OneDrive content.

CSA Resource Alignment

This campaign extends a lure pattern CSA has already documented in detail. The MacSync research note describes near-identical AI-brand impersonation, Anthropic Claude Code among the spoofed products, used to deliver a ClickFix-based macOS infostealer between November 2025 and February 2026, and recommends replacing long-lived local credentials with short-lived, identity-provider-issued alternatives and treating developer endpoints as high-privilege infrastructure ("MacSync Infostealer: ClickFix Campaigns via Fake AI Installers," Cloud Security Alliance, 2026) [3]. The ACR Stealer activity described here shows the same lure template operating against Windows endpoints and Microsoft 365 rather than macOS developer tooling, reinforcing that this is now a cross-platform template rather than a technique specific to one operating system or malware family.

The central risk in this campaign, an attacker reusing a stolen, already-authenticated session token to access cloud services without needing a password or a fresh MFA challenge, is precisely the abuse pattern addressed in CSA's "Using Zero Trust to Counter Identity Spoofing and Abuse" [4]. That publication's guidance on deploying User Behavior Analytics and Identity Threat Detection and Response tooling to flag anomalous account activity, and its emphasis on short-lived, automatically rotating credentials over long-term static tokens, maps directly onto the detection and remediation priorities Microsoft outlines for this campaign. CSA's AI Controls Matrix (AICM v1.1) provides the underlying control framework for organizations formalizing these practices, particularly within its Identity and Access Management domain, which governs credential scoping, session lifecycle management, and token revocation procedures of the kind this campaign's remediation guidance depends on ("AI Controls Matrix (AICM)," Cloud Security Alliance) [5].

References

[1] The Hacker News. "ACR Stealer Uses ClickFix Lures to Steal Browser Tokens and Microsoft 365 Files." The Hacker News, July 2026.

[2] Microsoft Threat Intelligence. "ACR Stealer: Two Observed Intrusion Chains Amid Increased Threat Activity." Microsoft Security Blog, July 16, 2026.

[3] Cloud Security Alliance. "MacSync Infostealer: ClickFix Campaigns via Fake AI Installers." CSA AI Safety Initiative, 2026.

[4] Cloud Security Alliance. "Using Zero Trust to Counter Identity Spoofing and Abuse." Cloud Security Alliance, 2026.

[5] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.