MacSync Infostealer: ClickFix Campaigns via Fake AI Installers

Social Engineering Bypasses macOS Security Controls to Harvest Enterprise Credentials

Unofficial AI-assisted Research

Cloud Security Alliance AI Safety Initiative

2026-03-16

Key Takeaways

Three documented ClickFix campaigns between November 2025 and February 2026 distributed the MacSync macOS infostealer by impersonating popular AI tools including OpenAI Atlas, ChatGPT, and Anthropic Claude Code [1]. Sophos researchers, who published primary analysis in March 2026, attribute the campaigns to infrastructure spanning 250 or more compromised WordPress sites across 12 countries [1]. Jamf Threat Labs independently flagged MacSync ClickFix distribution in December 2025, and Microsoft Defender findings from February 2026 corroborate coordinated deployment of MacSync alongside other infostealer families [3, 8].

The ClickFix technique exploits no software vulnerability. It succeeds entirely through social engineering: victims are manipulated into copying and pasting malicious terminal commands they believe are legitimate installation instructions [1, 2]. This approach bypasses Apple's Gatekeeper and notarization protections, which validate software packages rather than commands typed into the Terminal application. A later MacSync variant distributed through a code-signed and notarized installer further demonstrated the infostealer's evolving capacity to defeat macOS-native security controls [3]. Together, these developments signal that macOS endpoint security postures built primarily around Gatekeeper and signature verification are insufficient against contemporary infostealer tradecraft.

MacSync targets browser credentials, macOS Keychain contents, SSH keys, AWS credentials, Kubernetes configuration files, cryptocurrency wallets, and password manager vaults [2, 3, 4]. In enterprise environments, a single compromised developer or IT administrator workstation can expose cloud infrastructure access, CI/CD pipeline secrets, and production database credentials—providing initial access of the type sought by ransomware operators and data extortion groups. Organizations with significant macOS deployment, particularly those that issue developer workstations on Apple hardware, should treat this threat family as high priority.

Background

The MacSync Infostealer

MacSync is a macOS-targeting infostealer that emerged in April 2025 under the name Mac.c Stealer and was subsequently rebranded with enhanced capabilities [3]. It operates on a Malware-as-a-Service (MaaS) model, meaning it is available to multiple threat actors who pay for access to the malware's infrastructure and tooling. Jamf Threat Labs noted that MacSync "started appearing in detections in mid-2025 but infected hundreds of machines relatively fast" [9]—a pace consistent with rapid operator adoption following commercial availability of the MaaS platform.

The malware's core capability is the collection and exfiltration of credentials and sensitive files from compromised macOS systems. It harvests authentication material from Chromium- and Firefox-based browsers—including saved passwords, cookies, and browsing history—and directly reads the macOS Keychain database, which stores Wi-Fi passwords, application credentials, and certificate private keys [2, 4]. MacSync also targets the SSH key store at `~/.ssh`, AWS credential files at `~/.aws`, and Kubernetes configuration at `~/.kube`, making it particularly dangerous on developer workstations where these files routinely contain production-level access credentials [2]. Beyond credentials, the malware harvests documents matching extensions including PDF, DOCX, and spreadsheet formats, as well as data from Apple Notes and Telegram [4, 6].

Later MacSync variants extended the malware's capability in two directions. SHub Stealer, identified by Datadog Security Labs as an evolution of MacSync, added enterprise-oriented file collection—targeting CSV, XLS, XLSX, and RDP files limited to 5 MB—and implemented persistence through a fake GoogleUpdate binary registered as a LaunchAgent [4]. The variant also validates harvested macOS passwords using Directory Services before exfiltration, ensuring the credentials it transmits are functional [4]. A separate code-signed variant distributed through a fake video calling application installer demonstrated that threat actors had obtained an Apple Developer certificate and used it to create a notarized DMG that bypassed Gatekeeper without displaying any security warning to the victim [3].

The ClickFix Social Engineering Technique

ClickFix refers to a delivery technique, not a specific malware family. It deceives users into executing malicious shell commands by embedding those commands within pages or prompts styled to resemble legitimate technical content—installation guides, CAPTCHA verification flows, software error messages, or AI-generated documentation [1, 12]. The victim is instructed to open Terminal, paste a command, and

execute it. The command, which commonly takes the form of a `curl` pipe to `bash` or `zsh`, downloads and executes a payload without writing a recognizable installer to disk [1, 4]. Detections of ClickFix surged more than 500% between 2024 and 2025 [12], and it is now widely characterized as one of the fastest-growing social engineering techniques targeting enterprise environments.

Technical users, including developers and IT administrators, face elevated risk from ClickFix because `curl | sh` is a genuine and widely recognized installation pattern. Tools including Homebrew, many Python packages, and numerous cloud CLIs publish legitimate one-liner installation commands of this form. This cultural normalization raises the risk that even security-aware users may execute a ClickFix payload if the surrounding lure is sufficiently credible [1]. Importantly for organizations relying on phishing-resistant MFA, ClickFix is not addressed by FIDO2 authentication, which protects credential entry fields but has no bearing on terminal command execution [1].

Why macOS Is a Growing Infostealer Target

macOS represented approximately 27% of SME endpoints in 2025, up from 24% in 2024, and Apple devices account for approximately 15% of endpoints in large enterprise environments, with 73% of IT leaders reporting Apple fleet growth in the preceding year [11]. Palo Alto Unit 42 recorded a 101% increase in detected macOS infostealers between the last two quarters of 2024, with infostealers representing the largest category of new macOS malware in that year [10]. These figures reflect the expanding macOS attack surface; security researchers and vendors have attributed attacker interest in part to the perception—prevalent among macOS users—that Apple devices are inherently more secure than alternatives, a belief that may reduce vigilance [11].

A single compromised developer machine may yield cloud provider API keys, code repository tokens, CI/CD pipeline secrets stored in environment variable files, cryptocurrency wallet keys, and browser-stored credentials for SaaS platforms—the credential categories MacSync is specifically designed to harvest [4, 8]. These credentials frequently provide direct pathways into production infrastructure, making macOS developer endpoints a preferred initial access mechanism for financially motivated threat actors.

Security Analysis

Campaign Structure and Delivery Infrastructure

Sophos documented three distinct ClickFix campaigns distributing MacSync between November 2025 and February 2026 [1, 2]. The first, in November 2025, used Google sponsored search results to direct users searching for "ChatGPT Atlas" to fake landing pages hosted on `sites.google.com`, lending apparent legitimacy through the Google infrastructure domain before redirecting to a Bash script that fetched a MacSync binary [2]. The second campaign in December 2025 similarly exploited Google Ads, this time linking sponsored results to legitimate ChatGPT shared conversation pages—which themselves then redirected to malicious GitHub-themed landing pages [2]. The February 2026 campaign introduced more sophisticated evasion: dynamic AppleScript payloads executed in memory rather than written to disk, and infrastructure specifically checking whether a target system belonged to a corporate domain and what security tools were installed before proceeding [1].

Across these three campaigns, the threat actors used more than 250 compromised WordPress websites as distribution infrastructure, with confirmed victim geographies spanning Australia, Brazil, Canada, Czechia, Germany, India, Israel, Singapore, Slovakia, Switzerland, the United Kingdom, and the United States [1]. The geographic breadth and the investment in compromised legitimate infrastructure distinguish these campaigns from opportunistic commodity attacks and suggest organized operator activity.

Additional MacSync distribution activity documented by independent researchers adds further context. Moonlock Lab identified a Meta advertising campaign impersonating TradingView and OpenAI's Sora 2 video generator, directing victims to fraudulent download pages on domains including `tradingview.promo2026[.]com` and `sorachatgpts[.]com` [6]. HackRead documented a campaign that abused a Claude AI artifact titled "macOS Secure Command Execution" published on the official Anthropic platform; the artifact was viewed more than 15,600 times before detection and was promoted through hijacked verified Google Ads accounts belonging to a Canadian children's charity and a Colombian watch retailer [5]. Moonlock Lab separately identified 24 fake GitHub repositories, active since at least August 2025, impersonating more than 40 macOS applications across security, engineering, utility, and entertainment categories [7].

Evasion Techniques and Security Control Circumvention

MacSync and its associated delivery chain employ several techniques that specifically defeat macOS security controls. The ClickFix delivery mechanism requires no software vulnerability and bypasses code-signing validation entirely, because the malicious payload is fetched and executed by the legitimate system shell rather than installed as a discrete application [1, 2]. Gatekeeper validates application packages; it has no visibility into shell commands the user chooses to execute.

The code-signed and notarized MacSync variant documented by Jamf Threat Labs addressed even this distinction. The malware was distributed as a DMG file bearing a valid Apple Developer Team ID (`GNJLS3UYZ4`), which Apple has since revoked [3]. The DMG was inflated to 25.5 MB by embedding unrelated LibreOffice PDF documents, a known technique for reducing suspicion and evading file-size-based heuristics in sandbox analysis systems [3]. After installation, a helper executable retrieved an encoded script from the remote server and executed it, keeping the final payload off disk until execution time [3].

The SHub variant identified by Datadog implemented a rate-limiting mechanism—checking a file at `~/Library/Application Support/UserSyncWorker/last_up` and refusing to re-execute within a 3,600-second window—that reduces the forensic artifact density produced during analysis in automated sandbox environments [3, 4]. The February 2026 campaign variant introduced in-memory AppleScript payload execution, which evades static analysis tools that scan files on disk [1]. Threat actors have also implemented post-exfiltration self-destruction routines that delete malware artifacts after data transmission, complicating incident response [8].

Data Exfiltration and Enterprise Credential Exposure

MacSync bundles collected data into a ZIP archive—variously identified as `osalogging.zip` or `shub_log.zip` depending on the variant—and transmits it to attacker-controlled command-and-control servers [4, 5]. The credential categories targeted by MacSync align directly with the access paths most sought by ransomware operators and data extortion groups in enterprise environments. Browser-stored credentials in enterprise settings typically encompass SaaS platform logins, internal application credentials, and OAuth tokens for cloud services. Keychain contents on developer workstations frequently include database passwords, internal PKI certificates, and API credentials stored by integrated development environments. SSH keys and AWS credential files provide direct infrastructure access; in environments that rely on static credential authentication without behavioral monitoring, their theft can enable lateral movement without triggering alerting [2, 4].

Moonlock Lab's analysis of later MacSync variants documented targeting of more than 40 cryptocurrency wallets and more than 20 password manager applications [6]. The SHub variant specifically added enterprise file formats—including CSV, XLSX, and RDP connection files—to its collection scope, suggesting deliberate attacker adaptation to enterprise environments and an interest in selling or using harvested RDP credentials for further lateral movement [4].

Pillar Security, as cited in Sophos research [1], documented at least 20 distinct malware campaigns targeting AI and "vibe coding" tools between February and March 2026, indicating that the AI tool impersonation vector is not unique to the MacSync campaigns but represents a broader attacker strategy exploiting developer trust in this rapidly expanding tooling ecosystem.

Relationship to the Broader macOS Infostealer Ecosystem

MacSync operates within a documented and growing macOS infostealer market. Atomic macOS Stealer (AMOS), Banshee Stealer, Poseidon Stealer, Cthulhu Stealer, and DigitStealer are contemporaneous families active on the same distribution channels [10]. Microsoft Defender Security Research Team documented coordinated deployment of MacSync, DigitStealer, and AMOS in late 2025, with all three families targeting overlapping credential categories and using similar delivery techniques [8]. DigitStealer specifically implements country-of-origin checks—reading system locale settings and exiting on certain configurations—suggesting operator sophistication consistent with organized cybercriminal enterprises rather than opportunistic actors [8].

The ClickFix technique itself is not macOS-specific. Windows-targeting ClickFix campaigns distributing Lumma Stealer, Vidar, and other families were documented throughout 2024 and into 2025. The documented adaptation of ClickFix to macOS and its rapid proliferation across multiple threat actors reflects the technique's platform independence and low adaptation barrier—since ClickFix requires no exploit code and can be customized by modifying only the social engineering lure and the hosted payload—as well as the attractive credentials available on macOS developer workstations.

Recommendations

Immediate Actions

Organizations should implement endpoint detection rules for the specific behavioral patterns MacSync and related infostealers exhibit on macOS systems. At the filesystem level, rules should alert on creation of files at `/tmp/osalogging.zip`, `/tmp/shub_log.zip`, directories matching

`/tmp/shub_*`, and scripts at `/tmp/.c.sh`, as well as the addition of LaunchAgent plists impersonating Google services at `~/Library/LaunchAgents/com.google.keystone.agent.plist` [3, 4]. At the network level, detection should flag `curl` processes making connections to domains not in organizational allowlists, particularly where the curl output is piped directly to a shell interpreter. The parent-child process relationship `Terminal → curl → zsh` or `Terminal → curl → bash`, absent a preceding package manager invocation, should be treated as anomalous.

MacSync indicators of compromise available for immediate ingestion include the C2 domains `focusgroovy[.]com` and `gatemaden[.]space` documented by Jamf Threat Labs [3], and the Datadog-identified C2 infrastructure including `imper-strlk5[.]com`, `securityfenceandwelding[.]com`, `stobminipinporl[.]com`, `mini-zmoto[.]com`, and `mubasokurso[.]com` [4]. GitHub Pages domains used for ClickFix staging include `git-tool-install.github.io`, `io-app-git.github.io`, and `quadency-pro.github.io` [4]. The revoked Apple Developer Team ID `GNJLS3UYZ4` should be blocked in MDM policies where code signing validation is enforced [3].

Security teams should immediately review Google Ads account access for any organization whose Google Ads credentials are stored on developer endpoints. The HackRead-documented campaign demonstrated that attackers specifically sought and used verified Google Ads accounts to lend authenticity to malicious sponsored search results [5]; a compromised developer workstation yielding Google Ads credentials represents not only a data theft but a potential distribution infrastructure compromise.

Short-Term Mitigations

Endpoint management policy for macOS devices should be reviewed against the specific capabilities MacSync exploits. Apple's System Integrity Protection (SIP) should be confirmed enabled on all managed endpoints; while SIP does not prevent user-space persistence mechanisms such as LaunchAgent registration—the technique used by the SHub variant—it restricts modification of system directories and prevents kernel-level rootkit installation, reducing the severity of post-compromise scenarios. Organizations using Jamf Pro, Kandji, or equivalent MDM solutions should deploy configuration profiles that prevent user-space LaunchAgent installation from paths outside of system-managed directories, which would impede the persistence mechanism identified in the SHub variant [4].

Credential hygiene for SSH keys, AWS credentials, and Kubernetes configurations on developer workstations warrants review. Where possible, long-lived credential files stored at `~/.ssh`, `~/.aws`, and `~/.kube` should be replaced with short-lived credentials issued through identity-provider-

integrated tooling such as AWS IAM Identity Center, certificate-based SSH with a short validity period, or cloud-native workload identity mechanisms. Credentials that cannot be rotated to short-lived alternatives should be considered potentially compromised and rotated manually if any MacSync indicator of compromise is identified on an endpoint in the environment.

User awareness training for macOS users—and particularly for developers—should explicitly address the ClickFix technique. The training message should emphasize that unsolicited requests to paste commands into Terminal—appearing on public websites, in CAPTCHA flows, or in AI-generated documentation not originated from internal tooling—should be treated as high-suspicion and verified through a separate channel before execution [12].

Strategic Considerations

The documented abuse of Claude AI artifacts, GitHub Pages, Google Sites, and Medium as distribution infrastructure highlights a strategic gap in organizational security models: blocking known-malicious domains while permitting broad access to legitimate platform domains provides insufficient protection when attackers host malicious content on trusted platforms [5, 7]. Organizations should evaluate whether content security controls can be extended to monitor or restrict specific categories of user behavior—such as executing commands from browser-copied content—rather than relying solely on domain-based filtering.

Developer endpoint security programs should treat macOS developer workstations as equivalent to Tier 0 infrastructure in terms of monitoring and access control. The credential inventory available on a typical developer Mac—cloud API keys, repository tokens, Kubernetes access, SSH infrastructure access—frequently exceeds the credential value available on general employee endpoints and is of the same class as domain controller or production server credentials. Endpoint Detection and Response (EDR) solutions supporting macOS should be deployed and tuned for infostealer-specific behaviors, including anomalous `curl` invocations, Keychain database reads by non-system processes, and mass file collection from user home directories.

The MaaS model underlying MacSync means that disruption of any single operator does not remove the threat. As long as the underlying MacSync infrastructure remains available for purchase, new operators can acquire access and adapt delivery campaigns. Organizational resilience against this threat class requires architectural controls—short-lived credentials, MDM-enforced endpoint configuration, and network egress monitoring—that remain effective regardless of which campaign variant is in active circulation.

CSA Resource Alignment

The MacSync ClickFix campaigns map directly to the threat model described in CSA's MAESTRO framework for agentic AI systems. MAESTRO Layer 1 (Foundation Models and Execution Environment) and Layer 7 (Ecosystem and Infrastructure) are specifically implicated: the abuse of trusted AI platform content (Claude artifacts, ChatGPT-linked pages) as delivery infrastructure represents an attack on the trust model users place in AI service ecosystems. Organizations applying MAESTRO should extend their agentic AI threat models to include the developer workstation as an attack surface for credential exfiltration, recognizing that AI tool installation workflows are now an active initial access vector.

The CSA Cloud Controls Matrix (CCM) provides relevant control mappings for the mitigations described in this note. CCM Domain TVM (Threat and Vulnerability Management) addresses the need for endpoint-level detection controls, including behavioral anomaly detection for credential-access patterns that MacSync exhibits. CCM Domain IAM (Identity and Access Management) applies to the short-lived credential recommendations: IAM-04 (Credential Lifecycle and Provisioning Management) and IAM-07 (User Access Restriction/Authorization) support the replacement of long-lived filesystem-stored credentials with identity-provider-issued ephemeral access. CCM Domain SEF (Security Incident Management, E-Discovery, and Cloud Forensics) is relevant to the incident response complexities introduced by MacSync's self-deletion behavior.

CSA's Zero Trust guidance is directly applicable to the post-compromise scenario that MacSync enables. The core Zero Trust principle of "never trust, always verify" for network access—specifically, requiring continuous authentication validation rather than implicit trust based on network location or session continuity—limits the blast radius when SSH keys or cloud credentials are stolen. Organizations that have implemented Zero Trust network access (ZTNA) for cloud infrastructure reduce the exploitability of harvested credentials, since stolen keys must still satisfy device posture checks and identity provider validation before access is granted.

The CSA AI Organizational Responsibilities guidance supports the strategic recommendation regarding developer endpoint classification. Treating developer workstations as high-privilege assets subject to the same controls as production infrastructure is consistent with the organizational responsibility framework's emphasis on mapping AI system access to appropriate privilege tiers and ensuring that the endpoint security posture reflects the actual credential value present on the device.

References

[1] Jagadeesh Chandraiah, Tonmoy Jitu, Dmitry Samosseiko, Matt Wixey (Sophos X-Ops), "Evil evolution: ClickFix and macOS infostealers," Sophos, March 2026. https://www.sophos.com/en-us/blog/evil-evolution-clickfix-and-macos-infostealers

[2] Ravie Lakshmanan, "ClickFix Campaigns Spread MacSync macOS Infostealer via Fake AI Tool Installers," The Hacker News, March 16, 2026. https://thehackernews.com/2026/03/clickfix-campaigns-spread-macsync-macos.html

[3] Thijs Xhaflaire (Jamf Threat Labs), "From ClickFix to code signed: the quiet shift of MacSync Stealer malware," Jamf, December 22, 2025. https://www.jamf.com/blog/macsync-stealer-evolution-code-signed-swift-malware-analysis/

[4] Datadog Security Labs, "Tech impersonators: ClickFix and macOS infostealers," Datadog, February 10, 2026. https://securitylabs.datadoghq.com/articles/tech-impersonators-clickfix-and-macos-infostealers/

[5] Deeba Ahmed, "Google Ads and Claude AI Abused to Spread MacSync Malware via ClickFix," HackRead, February 16, 2026. https://hackread.com/google-ads-claude-ai-macsync-malware-clickfix/

[6] Ray Fernandez (Moonlock Lab), "ClickFix malware is infecting Macs via Facebook ads for fake AI apps," Moonlock, November 21, 2025. https://moonlock.com/clickfix-targeting-macs

[7] Ray Fernandez (Moonlock Lab), "MacSync Stealer Returns: SEO Poisoning and Fake GitHub Repositories," Moonlock, January 23, 2026. https://moonlock.com/macsync-stealer-fake-github-repositories

[8] Microsoft Defender Security Research Team, "Infostealers without borders: macOS, Python stealers, and platform abuse," Microsoft Security Blog, February 2, 2026. https://www.microsoft.com/en-us/security/blog/2026/02/02/infostealers-without-borders-macos-python-stealers-and-platform-abuse/

[9] Jamf Threat Labs, "Jamf Threat Labs dissects infostealer malware," Jamf, 2025. https://www.jamf.com/blog/infostealers-pose-threat-to-macos/

[10] Palo Alto Unit 42, "Stealers on the Rise: A Closer Look at a Growing macOS Threat," Palo Alto Networks, February 4, 2025. https://unit42.paloaltonetworks.com/macos-stealers-growing/

[11] KELA Cyber, "Watch Out: macOS Infostealers Are Rising," KELA, 2025. https://www.kelacyber.com/blog/the-rise-of-macos-infostealers-tactics-trends-defenses/

[12] Andrew Orr, "Bad CAPTCHA in the wild tricks Mac users into installing malware through Terminal," Apple Insider, March 10, 2026. https://appleinsider.com/articles/26/03/10/bad-captcha-in-the-wild-tricks-mac-users-into-installing-malware-through-terminal