ADGM THREAT INTELLIGENCE NEWSLETTER

PREPARED BY ADGM CTI

CATEGORY	ACTIONABLE
AUDIENCE	ADGM FSRA ENTITIES
DATE	09/7/2026
OVERALL THREAT SCORE	GUARDED
TARGET SECTOR	FINANCIAL SERVICES
TARGET REGION	MENA & GLOBAL
ATTRIBUTION	MULTIPLE
TLP	CLEAR

WEEKLY SUMMARY REPORT – 09 July 2026

14	5	0	0
Campaigns	**Vulnerabilities**	**Cyber Breach**	**Threat Actors**
Threat Campaigns of Potential Relevance to Financial Sector	Actively Exploited & Critical Vulnerabilities	Major Compromises and Breaches	Threat actor activities in the Middle East impacting Financial Sector

Summary

This week's cybersecurity newsletter highlights a rapidly evolving threat landscape where attackers increasingly exploit trusted technologies, AI ecosystems, remote management platforms, cloud services and software supply chains to gain access, steal credentials, enable ransomware, espionage and fraud operations. Key campaigns included AI-themed malware distribution, device-code phishing targeting Microsoft 365, abuse of autonomous AI agents, OAuth token theft, large-scale browser and software impersonation campaigns, supply-chain driven ransomware activity and sophisticated credential harvesting operations. The reporting also revealed growing risks from AI-centric attack techniques such as indirect prompt injection, phantom squatting, malicious AI-agent supply chains and AI-branded social engineering lures. Additionally, researchers documented active ransomware activity leveraging remote management tools, VPN access and Citrix infrastructure, alongside espionage-focused operations targeting government and regulated sectors. For financial institutions, these developments reinforce the need to prioritize credential security, phishing resilience, software supply-chain assurance, remote access monitoring and governance of AI-enabled technologies. Organizations should accelerate remediation of critical vulnerabilities affecting NetScaler, Langflow, SimpleHelp, Adobe Campaign Classic, and Apache Tomcat, while strengthening controls around AI agents, browser extensions, OAuth integrations, and third-party software. Continuous monitoring for credential theft, unauthorized remote access, token abuse and anomalous cloud activity remains essential to reducing exposure to these increasingly interconnected threats.

ADGM THREAT INTELLIGENCE SUMMARY

Ousaban Banking Trojan Campaign Uses Geofencing and Steganography [Campaign] [High]

Silent Swap Campaign Uses Malicious Browser Extension to Hijack Cryptocurrency Transactions [Campaign] [High]

Fake AI-Agent Threats Drive Infostealer Delivery, Brand Impersonation, and Autonomous Intrusion Activity [Campaign] [High]

ToddyCat Leverages OAuth Authorization Workflow for Email Access [Campaign] [High]

ARToken Expands Device Code Phishing Operations Against Microsoft 365 [Campaign] [High]

Anubis Ransomware Affiliates Abuse CitrixBleed 2 and Legitimate Remote Access Tools [Campaign] [High]

Indirect Prompt Injection Campaigns Target AI-Driven Workflows [Campaign] [Medium]

New BusySnake Stealer Enhances Armored Likho Espionage Campaign [Campaign] [Medium]

AI-Themed Browser Extension Enables Search Interception and Data Collection [Campaign] [Medium]

Credential Theft Enables Ransomware Deployment Through Vect and TeamPCP Collaboration [Campaign] [Medium]

ClickFix Campaign Uses Fake Google and Cloudflare Verification Lures [Campaign] [Medium]

Phantom Squatting Emerges as an AI Hallucination Driven Supply Chain Risk [Campaign] [Medium]

Large-Scale ScreenConnect Impersonation Campaign Distributing AsyncRAT [Campaign] [Medium]

Citrix Addresses Multiple High-Severity NetScaler Vulnerabilities [Vulnerability] [High]

Critical Vulnerabilities in Langflow OSS Could Lead to Full AI Platform Compromise [Vulnerability] [High]

SimpleHelp Authentication Bypass Vulnerability Actively Exploited [Vulnerability] [High]

Adobe Campaign Classic Security Update Addresses Remote Code Execution Risk [Vulnerability] [High]

Apache Tomcat Authentication Bypass Vulnerability [Vulnerability] [Medium]

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Ousaban Banking Trojan Campaign Uses Geofencing and Steganography	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified an ongoing 'Ousaban' banking Trojan campaign targeting users through phishing PDFs disguised as corrupted documents. The lure directs victims to a malicious website that performs geofencing and environmental validation before delivering a VBS (Visual Basic Script) downloader, which retrieves and executes the final Ousaban payload using steganography and multi-stage malware delivery techniques.

Organizations in the financial sector should be aware that the malware specifically monitors access to banking services and incorporates advanced evasion mechanisms, dynamic command-and-control resolution, and remote-control capabilities. These techniques could affect institutions by facilitating credential theft, unauthorized account activity, and subsequent malicious actions targeting online banking users.

Technical Details

- The campaign begins with a phishing PDF masquerading as a corrupted document and prompting users to click an update button that redirects them to a malicious webpage.
- The landing page performs geofencing checks and primarily targets users located in Spain and Portugal, restricting malware delivery to intended victims.
- Environmental validation includes analysis of language, time zone, IP-related information, and device characteristics to limit exposure and hinder analysis.
- In the latest variant, environmental assessment is moved to the server side, making detection criteria more difficult for researchers to observe.
- Approved victims receive a VBS script containing numerous benign function calls intended to obscure malicious activity.
- The VBS downloader retrieves a steganographic image containing an embedded ZIP archive and extracts the final Ousaban payload from it.
- After execution, the malware establishes persistence through a registry run key and records installation-related information for tracking purposes.
- Ousaban decrypts bank-related strings and monitors victim interaction with targeted banking services through web browsers.
- Rather than relying on a static configuration, the malware generates daily changing hostnames to resolve command-and-control infrastructure dynamically.
- Command-and-control functionality supports system profiling, screenshot capture, remote control, clipboard manipulation, keylogging, and deceptive victim-facing messages.

Recommendations

- Strengthen email security controls and inspect PDF attachments that contain embedded scripts or suspicious update prompts.
- Monitor and investigate execution of VBS scripts and other scripting activity originating from user download locations.
- Implement network monitoring to identify unusual outbound communications and dynamic infrastructure resolution patterns.
- Deploy endpoint detection capabilities to identify persistence mechanisms, remote-control activity, and unauthorized screenshot collection.
- Conduct user awareness training focused on phishing lures involving document updates, tax-related content, and fraudulent file access requests.

For the indicators of compromise (IOCs), refer to the attached CSV sheet. 

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Silent Swap Campaign Uses Malicious Browser Extension to Hijack Cryptocurrency Transactions	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified an active cryptocurrency-focused campaign that delivers a malicious browser extension disguised as a legitimate note-taking utility. The threat is distributed through unsigned installers that silently modify Chromium-based browser settings, deploy the extension without the standard installation process, and monitor clipboard activity to replace copied cryptocurrency wallet addresses with attacker-controlled alternatives during transactions.

Organizations in the financial sector should be aware that the campaign leverages browser trust abuse, dynamic wallet substitution, and blockchain-based infrastructure resolution to evade detection and maintain operational resilience. These techniques could affect cryptocurrency platforms, virtual asset providers and related services by facilitating unauthorized fund transfers and reducing opportunities for users to identify fraudulent transaction details before submission.

Technical Details

- The campaign deploys malicious Chromium browser extensions through unsigned ".NET" and Golang-based installers that masquerade the payload as a benign "Google Notes" application.
- The installer modifies protected browser preference files and recalculates integrity verification values, enabling the extension to appear legitimately installed.

- The extension requests extensive permissions, including access to all websites, browsing history, and clipboard content, providing broad visibility into user activity.
- Clipboard monitoring functions continuously inspect copied content and identify cryptocurrency wallet addresses using cryptocurrency-specific matching logic.
- When a wallet address is detected, the extension transmits the intercepted value to attacker-controlled infrastructure and retrieves a replacement address.
- The original wallet address is replaced prior to paste operations, causing cryptocurrency transfers to be redirected without the victim's awareness.
- The malware does not contain a static command-and-control location and instead queries blockchain infrastructure to dynamically obtain active backend information.
- Command-and-control details are decoded at runtime from blockchain query responses, allowing infrastructure changes without modifying the malware.
- The extension maintains persistence through browser configuration manipulation and can continue operating using locally cached backend information.
- Additional evasion mechanisms include installer self-deletion, deceptive application branding, dynamic wallet assignment, and selective exclusion of certain cryptocurrency-related websites from interception activities.

Recommendations

- Restrict installation and execution of unsigned software and browser extensions across enterprise-managed systems.
- Continuously audit installed browser extensions and investigate extensions requesting excessive access to websites, clipboard data, or browsing history.
- Monitor browser configuration files for unauthorized modifications that may indicate extension sideloading activity.
- Require users handling cryptocurrency transactions to independently validate destination wallet addresses before confirming transfers.
- Enhance endpoint and network monitoring to identify unauthorized browser modifications, clipboard manipulation activity, and suspicious outbound communications.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Fake AI-Agent Threats Drive Infostealer Delivery, Brand Impersonation and Autonomous Intrusion Activity	**HIGH**	**CLEAR**	**Campaign**	**CSC**

Executive Summary

Researchers have identified an active and expanding threat class that exploits trust in generative AI platforms through trojanized AI-tool installers, AI-vendor impersonation infrastructure, and the abuse of legitimate autonomous AI agents. Threat actors rapidly weaponize major AI product launches, leveraging malvertising, SEO (Search Engine Optimization) poisoning, AI-generated promotional content, and ClickFix social engineering to distribute Infostealers, credential-harvesting lures, and AI-enabled intrusion capabilities.

Organizations in the financial sector should be aware that confirmed exposure includes credential records linked to government, banking, telecommunications, aviation, healthcare, and critical-infrastructure entities. The convergence of credential theft, AI-brand impersonation, and autonomous agent abuse could affect institutions adopting AI-assisted workflows while increasing opportunities for account compromise, fraud, and unauthorized access to sensitive systems.

Technical Details

- The threat class consists of three primary vectors: trojanized AI-tool installers, AI-vendor impersonation infrastructure, and the abuse of legitimate autonomous AI agents for intrusion operations.
- Threat actors rapidly operationalize major AI product releases, using sponsored-search malvertising, SEO-poisoned download pages, AI-generated promotional videos, fake repositories, and ClickFix techniques to deliver malware.
- Trojanized AI-tool installers distribute commodity malware families including Amatera, AMOS, Vidar, Rhadamanthys, Lumma, GhostSocks, NOODLOPHILE, DinDoor, and Chaos rather than custom malware payloads.
- Fake AI download portals and documentation clones imitate legitimate vendor branding and are used to harvest credentials, payment information, session tokens, and other sensitive data.
- Some phishing infrastructure employs cloaking mechanisms and allowlist-based visitor filtering to restrict visibility and reduce detection by researchers and automated scanners.
- ClickFix social engineering campaigns instruct victims to paste attacker-supplied commands into terminal or system execution dialogs, turning users into the execution mechanism.
- Infostealer payloads collect browser credentials, cookies, session tokens, autofill information, cryptocurrency-wallet data, and other sensitive information before exfiltration.
- Command-and-control resilience is enhanced through techniques such as blockchain-based configuration retrieval and the use of resilient communication channels for stolen data distribution.
- Autonomous AI-agent abuse enables legitimate coding agents and connected tool servers to perform reconnaissance, vulnerability discovery, credential collection, and data exfiltration with limited operator intervention.

- Additional risks arise from malicious AI marketplace skills, compromised tool servers, hallucinated dependency packages, and prompt-injection attacks targeting agent workflows and supply chains.

Recommendations

- Restrict software downloads to approved sources and block access to known AI-themed impersonation sites and unauthorized repositories.
- Prioritize credential resets, session invalidation, and exposure validation for users potentially impacted by infostealer infections.
- Monitor for ClickFix-style activity, including suspicious command execution initiated through user interaction with websites or social engineering lures.
- Treat AI agents, connected tools, marketplace apps, and third-party software as potential security risks, and regularly monitor, review, and assess them for security issues.
- Enhance detection and hunting for infostealer behavior, credential harvesting, anomalous authentication activity, and unauthorized data collection across enterprise environments.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure – Fake AI-Agent Threats Drive Infostealer Delivery, Brand Impersonation, and Autonomous Intrusion Activity

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
ToddyCat Leverages OAuth Authorization Workflow for Email Access	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified a ToddyCat campaign targeting corporate Gmail accounts by abusing the OAuth authorization workflow and active browser sessions. The threat actor deploys a custom tool named Umbrij through DLL sideloading, connects to Chromium-based browsers via remote debugging ports, and obtains OAuth authorization tokens that provide API-based access to email resources without requiring direct credential theft.

Organizations in the financial sector should be aware that this technique focuses on accessing business communications while avoiding traditional security monitoring mechanisms. The campaign could affect organizations that rely on cloud-hosted email services by enabling unauthorized access to corporate correspondence through trusted browser sessions and legitimate authentication workflows.

Technical Details

- The campaign targets corporate Gmail accounts and leverages the OAuth 2.0 authorization framework to obtain access tokens that enable access to mailbox resources through Google APIs.

- Threat actors developed a custom ".NET"- based tool named 'Umbrij' that automates the attack workflow and is designed to operate on Chromium-based browsers.
- Initial execution was observed through malicious scheduled-task activity combined with DLL sideloading techniques that leveraged legitimately signed applications.
- The malware uses legitimate executables vulnerable to DLL sideloading to load and execute the 'Umbrij' payload while reducing suspicion.
- Before launching the attack, 'Umbrij' verifies port availability, prepares the environment, and retrieves the user context required for browser interaction.
- The tool duplicates user access tokens from active explorer.exe processes, allowing execution within the targeted user's security context.
- 'Umbrij' enumerates browser profiles and identifies accounts associated with authenticated Google services by examining profile information stored by the browser.
- The malware creates backup directories and copies browser artifacts, including profile data, storage repositories, credentials, session-related files, browser preferences, and autofill information.
- The tool launches Chrome or Edge in headless mode with remote debugging enabled, using copied profile data to inherit existing authenticated sessions.
- Through browser automation and remote debugging access, 'Umbrij' obtains OAuth authorization codes and exchanges them for access tokens, enabling unauthorized access to cloud-hosted email resources.

Recommendations

- Monitor for unusual browser launches utilizing headless mode, remote debugging parameters, or non-standard browser profile locations.
- Investigate DLL sideloading activity involving trusted executables and unexpected scheduled tasks on corporate endpoints.
- Review OAuth application activity and monitor for unauthorized token generation or unusual API access patterns.
- Implement enhanced monitoring for browser-profile access, credential stores, and large-scale copying of browser-related files.
- Conduct regular reviews of cloud email access logs and investigate abnormal authentication activity originating from trusted user sessions.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure – ToddyCat Leverages OAuth Authorization Workflow for Email Access

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
ARToken Expands Device Code Phishing Operations Against Microsoft 365	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers at Microsoft has identified ARToken, a phishing-as-a-service platform targeting Microsoft 365 accounts through device code phishing attacks. The operation uses vendor-impersonation invoice lures, malicious SharePoint-themed links, and OAuth device authorization workflows to capture authentication tokens, enabling account access while bypassing traditional credential theft and multi-factor authentication challenges.

Organizations in the financial sector should be aware that ARToken provides affiliates with a comprehensive post-compromise toolkit supporting token persistence, email access, business email compromise activity, and document exfiltration. These capabilities could affect organizations that rely heavily on Microsoft 365 by facilitating long-term access to corporate communications, cloud resources, and sensitive business information.

Technical Details

- ARToken operates as a multi-tenant phishing-as-a-service platform and shares infrastructure, operational patterns, and OAuth-based token capture mechanisms with previously documented device code phishing operations.
- Initial access relies on targeted vendor-impersonation emails that leverage existing business relationships and use invoice-themed lures to increase the likelihood of user interaction.
- Email messages contain links that appear to reference legitimate SharePoint resources while redirecting victims to attacker-controlled Microsoft 365 environments.
- The phishing workflow abuses Microsoft's OAuth device authorization process, generating device codes that direct victims to complete authentication through legitimate Microsoft login pages.
- Once loaded, the phishing kit extracts victim information, requests device authorization data, and presents time-limited authentication instructions to targeted users.
- The platform incorporates a seven-layer anti-analysis framework that evaluates browser characteristics, user interaction patterns, automation indicators, timing behavior, and movement telemetry before exposing payloads.
- JavaScript payloads are encrypted and decrypted at runtime, complicating static analysis, and reducing visibility for automated scanning tools.
- Captured authentication tokens can be refreshed, exported, shared between operators, and elevated into more persistent access mechanisms that extend account access longevity.
- The platform provides direct access to compromised Outlook mailboxes, including email review, message transmission, attachment access, inbox-rule manipulation, and keyword-based monitoring.

- Additional functionality supports SharePoint and OneDrive access, document management, phishing infrastructure deployment, and collaborative management of compromised accounts through a centralized operator dashboard.

Recommendations

- Monitor for device code authentication activity and investigate unexpected device authorization requests involving corporate Microsoft 365 accounts.
- Implement conditional access controls and strengthen monitoring of OAuth-based authentication events and token-related activity.
- Review email security controls for invoice-themed phishing attempts, vendor impersonation, and suspicious SharePoint links.
- Continuously monitor Microsoft 365 environments for unauthorized inbox rules, abnormal mailbox access, and unusual file activity within SharePoint and OneDrive.
- Conduct user awareness training focused on device code phishing techniques, unsolicited authentication prompts, and business-themed social-engineering campaigns.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure - ARToken Expands Device Code Phishing Operations Against Microsoft 365

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Anubis Ransomware Affiliates Abuse CitrixBleed 2 and Legitimate Remote Access Tools	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified multiple Anubis ransomware intrusions leveraging both exploitation of the "CitrixBleed 2" vulnerability and the use of valid VPN credentials to gain initial access. Following access, affiliates blend malicious activity with legitimate administration by deploying remote management tools, conducting credential theft, moving laterally through enterprise networks, and establishing alternate access channels before ultimately deploying ransomware.

Organizations in the financial sector should be aware that the observed tradecraft focuses on high-value infrastructure, including remote access services, domain controllers, hypervisors, backup systems, and storage platforms. These techniques could affect organizations by enabling persistent unauthorized access, large-scale credential compromise, data exfiltration, and operational disruption through coordinated ransomware deployment.

Technical Details

- Initial access was obtained through exploitation of "CitrixBleed 2", a pre-authentication memory disclosure vulnerability, as well as through the use of valid VPN credentials originating from hosting-provider infrastructure.
- Following VPN (Virtual Private Network) access, threat actors conducted authentication activity involving RDP (Remote Desktop Protocol) and SMB (Server Message Block), enabling broader access to enterprise systems and supporting subsequent attack stages.
- RDP was extensively used for hands-on-keyboard operations, with affiliates pivoting between servers, remote desktop infrastructure, file servers, domain controllers, backup platforms, and hypervisors.
- PsExec (Process Execute) service creation was repeatedly observed alongside lateral movement activity, remote administration operations, and ransomware staging efforts.
- Affiliates deployed legitimate remote management tools including ScreenConnect, Zoho Assist, MeshAgent, Remotely, UltraVNC, mRemoteNG, and Total Software Deployment to maintain access while blending in with normal administrative activity.
- In several incidents, actors attempted to create alternative outbound access channels using Cloudflare Tunnel functionality, authenticated proxy services, and SSH-based SOCKS tunneling.
- Credential harvesting activity included the use of Mimikatz, browser credential exports, and extraction of Active Directory database information to expand access across compromised environments.
- Tools used to access and transfer data, such as S3 Browser, rclone, s5cmd, WinSCP, and PuTTY, were used along with activities aimed at stealing credentials and preparing for ransomware attacks.
- Defense-evasion actions included endpoint security tampering, security software removal attempts, log clearing, and deletion of ransomware-related artifacts after execution.
- The final attack stage involved deployment of Anubis ransomware across Windows and Linux systems, resulting in encrypted files and ransom-note creation throughout affected environments.

Recommendations

- Immediately patch internet-facing Citrix NetScaler systems and review active sessions for signs of unauthorized access.
- Monitor VPN authentication activity for logins originating from hosting providers, unusual geolocations, or abnormal user behavior.
- Establish strict controls around approved remote management tools and investigate unauthorized deployments across enterprise assets.
- Implement detection and response mechanisms for credential-dumping activity, Active Directory database access, and lateral movement through RDP and PsExec.
- Monitor for unauthorized tunneling tools, cloud-transfer utilities, and outbound connectivity that may indicate persistence or data-exfiltration activity.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Indirect Prompt Injection Campaigns Target AI-Driven Workflows	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified multiple indirect prompt injection campaigns that target AI-driven workflows by embedding hidden instructions within malicious websites. The campaigns combine SEO (Search Engine Optimization) poisoning, structured metadata abuse, typosquatting, and concealed prompt content to influence AI agents that browse web content, potentially altering decision-making processes and directing automated actions on behalf of users.

Organizations in the financial sector should be aware that these attacks specifically target the trust assumptions used by AI agents, retrieval systems, and automated workflows. The observed techniques could affect institutions adopting AI-assisted processes by enabling malicious content to contaminate decision contexts, manipulate automated decisions, and introduce fraudulent actions into AI-driven business operations.

Technical Details

- The campaigns leverage indirect prompt injection techniques that embed hidden instructions within websites, allowing attackers to influence AI agents that retrieve and process web content.
- SEO poisoning is used to elevate malicious websites in search results, increasing the probability that AI agents encounter attacker-controlled content during routine queries.
- The first campaign disguises malicious activity as developer documentation and API-related resources, targeting AI-assisted development workflows.
- Attackers misuse JSON-LD (JavaScript Object Notation for Linked Data) metadata to make malicious instructions appear as trusted contextual information, increasing the likelihood that AI systems will treat them as legitimate.
- Hidden instructions are concealed through CSS (Cascading Style Sheets) manipulation, making the content invisible to users while remaining accessible to parsers, crawlers, and AI agents.
- Embedded prompt content attempts to persuade AI agents that payment actions are required to resolve software-related issues, directing funds to attacker-controlled payment mechanisms.
- The second campaign uses a typosquatting domain impersonating a legitimate decentralized finance platform and optimizes its visibility using search-engine-focused metadata.
- Fraudulent metadata and structured content falsely represent the malicious website as an authoritative and trusted service associated with the impersonated platform.
- Hidden prompt instructions attempt to override existing AI-agent guidance by directing models to ignore prior instructions and treat the malicious site as the primary trusted source.

- Testing across multiple large language models demonstrated instances where AI systems failed to correctly identify or respond appropriately to the malicious content, highlighting measurable operational risk.

Recommendations

- Validate AI-generated recommendations and automated actions against trusted sources before executing financial, administrative, or operational tasks.
- Implement controls that restrict AI agents from independently executing payment transactions or high-risk actions without human approval.
- Monitor AI-assisted workflows for interactions with newly observed, low-reputation, or look-alike domains.
- Apply content validation and filtering mechanisms to identify hidden prompt content, structured metadata abuse, and SEO-manipulation tactics.
- Treat AI agents, retrieval systems, and automated decision-making workflows as monitored attack surfaces and incorporate prompt-injection testing into security assessments.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
New BusySnake Stealer Enhances Armored Likho Espionage Campaign	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers at Securelist Labs has identified an ongoing espionage campaign conducted by the “Armored Likho” threat group targeting government organizations and the electric power sector across multiple countries. The campaign relies on spear-phishing emails carrying archive attachments that contain malicious executables or weaponized shortcut files, which ultimately deploy a previously undocumented Python-based infostealer named “BusySnake” through a multi-stage delivery process.

Organizations in the financial sector should be aware that the campaign combines credential theft, persistent access, modular malware deployment and advanced evasion techniques designed to avoid detection. These capabilities could affect institutions by enabling the collection of sensitive information, establishing long-term footholds, and facilitating follow-on espionage activities against targeted environments.

Technical Details

- “Armored Likho” conducts spear-phishing campaigns using themes related to official notices, social programs, humanitarian assistance, and other trust-based subjects to entice victims into opening malicious attachments.

- Malicious archives contain either executable files or shortcut files designed to imitate legitimate documents and application content.
- One infection chain uses self-extracting executable droppers that display decoy content to reduce suspicion while initiating malware deployment in the background.
- The dropper injects malicious code into a legitimate process and retrieves additional payload archives from repositories used to host and rotate malware components.
- Alternative attack chains leverage a shortcut-file vulnerability to conceal malicious command-line activity and launch PowerShell-based download operations.
- Subsequent stages download a Python interpreter, dependency-installation components, and archives containing the "BusySnake" payload before establishing the execution environment.
- The malware creates scripts that remove initial loader components and establishes persistence through scheduled tasks configured to repeatedly execute the payload.
- Researchers observed indications that some first-stage loader components were generated using large language models, including unusual coding patterns and extensive automated comments.
- "BusySnake" is a Python-based infostealer that uses code obfuscation and runtime decryption mechanisms to hinder static analysis and detection efforts.
- The malware operates through modular handlers and maintains configuration data that includes communication settings, collection rules, operational parameters, and command-and-control information.

Recommendations

- Strengthen email security controls to identify and block archive-based phishing attachments and suspicious shortcut files.
- Monitor for unauthorized PowerShell execution, code injection activity, and unexpected downloads of scripting runtimes or dependency-management tools.
- Investigate scheduled-task creation and recurring script executions that could indicate persistence mechanisms.
- Deploy behavioral monitoring capable of detecting obfuscated Python-based malware and unusual credential-collection activity.
- Conduct user awareness training focused on phishing emails using government, humanitarian, and administrative-themed lures.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
AI-Themed Browser Extension Enables Search Interception and Data Collection	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers at Microsoft Threat Intelligence has identified a malicious Chromium-based browser extension masquerading as an AI-powered search tool associated with Perplexity AI. The extension uses AI-related branding, a typo-squatted domain, and browser search-provider overrides to intercept search queries and real-time search suggestions, routing user activity through attacker-controlled infrastructure before redirecting traffic to legitimate search engines.

Organizations in the financial sector should be aware that the extension enables continuous collection of browsing activity and user search behavior while maintaining the appearance of normal search functionality. These capabilities could affect organizations by exposing sensitive business research, user interests, and potentially valuable operational information through unauthorized monitoring of browser activity.

Technical Details

- The malicious extension impersonates the Perplexity AI brand and uses a typo-squatted domain designed to create confusion regarding its legitimacy and affiliation.
- The extension was distributed through the Chromium ecosystem and presented itself as an AI-powered search utility to encourage installation by unsuspecting users.
- Analysis identified browser search-provider overrides that forcibly set the extension as the default search provider within Chromium-based browsers.
- All browser searches are routed through attacker-controlled infrastructure before users are redirected to expected search providers, allowing interception of search activity.
- The extension captures not only completed search queries but also real-time search suggestions, enabling visibility into characters typed into the browser address bar.
- Search suggestions and query traffic are transmitted to infrastructure not associated with the legitimate AI service, creating a central collection point for user activity.
- The extension abuses browser search settings to transparently redirect users while preserving the appearance of legitimate search-engine results.
- Advanced network-manipulation permissions are requested, allowing traffic redirection, request rewriting, interception, and monitoring of redirection activity.
- Researchers observed the use of declarative network request functionality to create a two-stage process in which attacker-controlled systems process search traffic before redirection occurs.
- The campaign demonstrates how threat actors continue to leverage AI-themed branding and trusted technology trends as social-engineering mechanisms to increase installation rates and reduce user suspicion.

Recommendations

- Restrict browser extension installations to approved repositories and enforce allowlisting for enterprise-managed browsers.
- Audit installed browser extensions for unauthorized search-provider modifications and excessive network-related permissions.
- Monitor for unexpected browser traffic routed through third-party infrastructure or non-approved search services.
- Educate users on AI-themed social-engineering tactics, brand impersonation, and the risks associated with installing unverified browser extensions.
- Regularly review browser configurations and remove extensions that request capabilities inconsistent with their advertised functionality.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Credential Theft Enables Ransomware Deployment Through Vect and TeamPCP Collaboration	**MEDIUM**	**CLEAR**	**Campaign**	**Open Source**

Executive Summary

Researchers have identified a coordinated campaign involving the Vect ransomware and TeamPCP, combining large-scale credential harvesting from supply chain compromises with ransomware deployment. The partnership leverages compromised software repositories, poisoned software updates, and credential-stealing implants introduced into widely used development tools, enabling access to numerous downstream organizations before ransomware execution.

Organizations in the financial sector should be aware that the campaign demonstrates how software supply chain breaches can be rapidly converted into large-scale ransomware operations. These activities could affect enterprises that rely on third-party software, development platforms, and cloud-based environments by exposing credentials, enabling unauthorized access, and facilitating follow-on extortion activities across interconnected business ecosystems.

Technical Details

- TeamPCP and Vect established a formal operational partnership, combining credential theft, data exfiltration, and ransomware deployment capabilities into a coordinated attack model.
- TeamPCP previously conducted large-scale exploitation campaigns targeting exposed development and cloud infrastructure, impacting organizations across multiple sectors and regions.

- A major phase of the operation involved compromising software development environments and injecting credential-harvesting malware into trusted software distribution channels.
- Malicious software updates were designed to continue performing expected application functions while secretly collecting passwords, cloud credentials, and other sensitive information from victim environments.
- Stolen credentials were subsequently used to conduct follow-on compromises, enabling attackers to access additional software repositories, development pipelines, and enterprise environments.
- Researchers observed the deployment of a self-propagating worm that spread through numerous software packages, increasing the reach of the compromise across the software ecosystem.
- Additional supply chain attacks targeted software development tools, plugins, automation workflows, and AI-related platforms used by large numbers of organizations.
- Some malicious packages incorporated credential-stealing functionality and mechanisms designed to execute automatically within development or runtime environments.
- The operation resulted in extensive credential collection, large-scale data exfiltration, and compromise of numerous enterprise software-as-a-service environments.
- The partnership enabled Vect affiliates to use credentials obtained through TeamPCP compromises to conduct ransomware deployments against affected organizations, establishing an operational path from supply chain compromise to ransomware execution.

Recommendations

- Maintain a comprehensive inventory of third-party software, development tools, plugins, and dependencies used across the organization.
- Implement strict validation and integrity verification procedures before deploying software updates into production environments.
- Monitor for unusual authentication activity involving development platforms, cloud services, and software repositories.
- Rotate credentials and review privileged access controls following any software supply chain incident affecting deployed technologies.
- Enhance monitoring for credential theft, unauthorized repository access, abnormal CI/CD pipeline activity, and indicators of ransomware preparation.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
ClickFix Campaign Uses Fake Google and Cloudflare Verification Lures	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers at Malwarebytes has identified an evolving ClickFix campaign that uses fake Google and Cloudflare verification pages to trick users into executing malicious PowerShell commands on their own systems. The attack leverages fraudulent human-verification prompts, Google-themed lures, and compromised or repurposed websites to persuade victims to manually execute attacker-supplied commands, resulting in the delivery of multiple malware families through a shared infrastructure.

Organizations in the financial sector should be aware that the campaign combines social engineering with malware-delivery mechanisms capable of stealing credentials, establishing remote access, and deploying additional payloads. These techniques could affect organizations by enabling account compromise, unauthorized system access, and broader intrusion activity following successful user interaction with the malicious verification pages.

Technical Details

- The campaign relies on ClickFix social engineering techniques that instruct users to manually copy and execute malicious commands under the guise of completing verification or troubleshooting activities.
- Multiple malware families have been distributed through the infrastructure, including HijackLoader, StealC, Remus, Amatera Stealer, CastleLoader, NetSupport, a Rust-based stealer, and a previously undocumented loader known as ResiLoader.
- One observed infection chain used a trojanized version of a legitimate messaging application that deployed ResiLoader, which subsequently disabled security controls before delivering an infostealer payload.
- The campaigns have been active since at least late 2025 and continue to evolve by testing new delivery methods, infrastructure, and payload combinations.
- Common characteristics include the use of PowerShell-based download-and-execute commands and staged malware delivery from attacker-controlled infrastructure.
- Payload distribution has been observed through repurposed legacy websites, compromised websites, Cloudflare Pages infrastructure, and fraudulent online services.
- Fake Google reCAPTCHA pages are used to present users with “Manual Verification Required” workflows that copy malicious commands to the victim’s clipboard.
- Additional verification lures use obfuscated HTML and custom code frameworks capable of retrieving malicious commands from local or remote sources.
- Some campaign variants incorporate an approval mechanism that allows attackers to decide which command a victim receives and executes.

- Recent activity includes Google Meet-themed lures that claim to resolve audio issues by instructing users to execute malicious commands, further broadening the range of social-engineering themes.

Recommendations

- Educate users that legitimate services such as Google, Cloudflare, and collaboration platforms will not require manual execution of PowerShell or terminal commands for verification.
- Block and investigate PowerShell execution initiated immediately after browser interactions or clipboard-based command delivery.
- Monitor for suspicious download-and-execute activity originating from user-initiated scripting engines.
- Enhance web filtering controls to restrict access to newly observed, compromised, or suspicious verification-related websites.
- Deploy endpoint monitoring capable of detecting loader activity, credential theft behavior, remote-access tools, and multi-stage malware delivery chains.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Phantom Squatting Emerges as an AI Hallucination Driven Supply Chain Risk	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified an emerging threat known as 'Phantom Squatting', where threat actors register web domains hallucinated by large language models and weaponize them for malicious purposes. By systematically probing AI systems for fictitious brand-related domains, attackers can pre-register these domains and position them as trusted destinations that may later be recommended by AI assistants, coding tools, and autonomous agents.

Organizations in the financial sector should be aware that this threat targets the growing trust placed in AI-generated recommendations and automated workflows. These techniques could affect organizations using AI-assisted development, research, and automation platforms by directing users or AI agents to attacker-controlled infrastructure that bypasses traditional phishing delivery methods and reputation-based security controls.

Technical Details

- Researchers found that large language models consistently generate non-existent domains associated with legitimate brands, creating opportunities for adversaries to register and weaponize those domains.
- The attack technique, termed phantom squatting, extends the concept of AI-generated supply chain risks from hallucinated software packages to hallucinated web infrastructure and service endpoints.

- Threat actors conduct adversarial probing of AI systems to identify recurring hallucinated domains associated with targeted organizations, brands, or services.
- Once identified, attackers rapidly register selected domains and have been observed deploying malicious content within hours of registration.
- Researchers documented a case where a phishing kit was deployed against a domain that had been identified as a high-risk hallucination target weeks before its registration.
- AI assistants may inadvertently function as delivery mechanisms by recommending hallucinated domains as legitimate websites, portals, APIs, or service endpoints.
- Autonomous AI agents and developers may subsequently interact with or integrate these attacker-controlled domains into operational workflows and software environments.
- Newly registered phantom domains benefit from a zero-reputation state, enabling them to evade reputation-based security controls that rely on historical intelligence and prior observations.
- Researchers analyzed hundreds of global brands and generated millions of AI-produced URLs, identifying thousands of malicious URLs and a substantial inventory of unregistered hallucinated domains that remain available for abuse.
- Key attacker advantages include cross-model consistency of hallucinated domains, persistence of domain generation across different AI configurations, and the ability to exploit trusted AI-generated outputs without relying on traditional phishing channels.

Recommendations

- Implement validation controls requiring independent verification of domains, URLs, and service endpoints generated by AI systems before operational use.
- Monitor AI-assisted development, automation, and research workflows for references to newly observed or previously unknown domains.
- Incorporate AI-generated infrastructure risks into third-party and supply chain security programs.
- Apply enhanced scrutiny to newly registered domains referenced in AI-generated content, particularly those associated with trusted brands or business services.
- Treat AI assistants, coding platforms, and autonomous agents as potential supply chain dependencies and continuously assess their outputs for malicious or fabricated recommendations.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Large-Scale ScreenConnect Impersonation Campaign Distributing AsyncRAT	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified a large-scale malware distribution campaign that uses spoofed software-download websites impersonating popular applications to deliver malicious 'ScreenConnect' installations and AsyncRAT payloads. Victims are lured into downloading trojanized installer archives that appear to provide legitimate software, but instead use DLL sideloading techniques to deploy a hidden remote administration service that grants attackers persistent access to compromised systems.

Organizations in the financial sector should be aware that the campaign combines trusted software impersonation, remote-access abuse, and multi-stage malware execution to establish long-term control over victim devices. These techniques could affect organizations by enabling unauthorized remote access, security-control bypass, malware deployment, and follow-on credential theft or espionage activity across enterprise environments.

Technical Details

- The campaign uses fraudulent websites impersonating widely used software products, including utilities, multimedia applications, and gaming-related tools, to distribute malicious installers.
- Researchers identified more than 90 spoofed domains localized across multiple languages, indicating a broad and scalable distribution operation.
- Malicious installer archives bundle a legitimate, digitally signed Microsoft executable together with a malicious library used for DLL sideloading.
- The sideloaded library installs and launches a hidden 'ScreenConnect' remote administration service, providing attackers with persistent remote access to compromised endpoints.
- Investigation of one compromise revealed 'ScreenConnect' being used to execute PowerShell and VBScript payloads that initiated subsequent attack stages.
- Malicious PowerShell activity modified Microsoft Defender settings, added extensive exclusions, and reduced security visibility across the affected host.
- The scripts also altered system settings to suppress User Account Control prompts, reducing barriers to further malicious execution.
- Subsequent VBScript activity generated multiple files and launched additional payloads designed to conceal activity and continue the infection chain.
- Later-stage PowerShell components decrypted and loaded malicious code in memory using encoded data, custom decryption routines, and reflective execution techniques.
- The malware ultimately leveraged process hollowing techniques to execute AsyncRAT, enabling remote control and additional malicious operations on compromised systems.

Recommendations

- Restrict software downloads to trusted vendor sources and actively block access to unauthorized software-distribution websites.
- Monitor for unexpected installation and execution of remote administration tools, particularly 'ScreenConnect' instances not approved by IT teams.
- Investigate PowerShell, VBScript, and DLL sideloading activity originating from newly installed software packages.
- Review endpoint security configurations for unauthorized Defender exclusions, security-policy changes, and UAC-related modifications.
- Deploy behavioral detection capabilities for process hollowing in-memory payload execution, reflective loading, and remote-access malware activity.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Citrix Addresses Multiple High-Severity NetScaler Vulnerabilities	HIGH	CLEAR	Vulnerability	CSC

Executive Summary

Citrix has disclosed six vulnerabilities affecting NetScaler ADC (Application Delivery Controller) and NetScaler Gateway that could allow unauthenticated attackers to read arbitrary files, trigger memory corruption, cause denial-of-service conditions, or disclose sensitive information depending on the enabled features and deployment configuration. The vulnerabilities affect widely deployed network access and application delivery infrastructure, including gateway, authentication, SAML (Security Assertion Markup Language) identity provider, DNS, and management interfaces, increasing the potential exposure of internet-facing environments.

Organizations in the financial sector should be aware that several of the vulnerabilities require only network access to exposed NetScaler services and do not require authentication. Successful exploitation could affect the confidentiality, availability, and security of critical remote-access infrastructure. Immediate patching and implementation of the vendor-recommended mitigations are advised, particularly for systems providing VPN, authentication, or application delivery services.

Technical Details

- Six vulnerabilities were disclosed affecting NetScaler ADC and NetScaler Gateway deployments.
- CVE-2026-8451 is an out-of-bounds memory read vulnerability caused by insufficient input validation and affects deployments configured as a SAML Identity Provider (IdP).

- CVE-2026-8452 is a memory overflow vulnerability that may lead to unpredictable behavior and denial-of-service conditions on appliances configured as Gateway services or AAA virtual servers.
- CVE-2026-8655 includes multiple memory overflow vulnerabilities affecting Oracle Load Balancer, DNS Proxy, and DNS Recursive Resolver deployments.
- CVE-2026-10816 is an unauthenticated arbitrary file read vulnerability that may allow access to files when management interfaces are reachable through NSIP, Cluster Management IP, or management-enabled SNIP interfaces.
- CVE-2026-10817 is an out-of-bounds memory read vulnerability affecting configurations where TCP Timestamp is enabled in associated TCP profiles.
- CVE-2026-13474 allows denial-of-service through malformed HTTP/2 requests when HTTP/2 is enabled in associated HTTP profiles.
- Three vulnerabilities (CVE-2026-8451, CVE-2026-8452, and CVE-2026-8655) have a CVSS v4.0 score of 8.8 (High).
- CVE-2026-13474 has a CVSS v4.0 score of 8.7 (High), while CVE-2026-10816 is rated 7.1 (High).

Recommendations

- Upgrade all affected NetScaler ADC and NetScaler Gateway appliances to the latest supported firmware.
- Apply the additional HTTP/2 configuration required to fully mitigate CVE-2026-13474.

Vulnerability and affected product details can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Critical Vulnerabilities in Langflow OSS Could Lead to Full AI Platform Compromise	**HIGH**	**CLEAR**	**Vulnerability**	**CSC**

Executive Summary

IBM Security researchers have disclosed six critical vulnerabilities affecting Langflow OSS (Open-Source Software), an open-source platform used to develop AI and large language model (LLM) applications. The vulnerabilities impact Langflow versions 1.0.0 through 1.10.0 and include weaknesses that may enable unauthenticated remote code execution, authorization bypass, insecure deserialization, command execution, cross-tenant access, and credential abuse. Several vulnerabilities have demonstrated proof-of-concept exploitation, significantly increasing the risk to exposed deployments.

Organizations in the financial sector should be aware that successful exploitation could result in complete compromise of AI application environments, unauthorized access to sensitive data, manipulation of AI workflows, and cross-tenant attacks. Given the critical severity of the vulnerabilities and the potential impact on AI-enabled business services, immediate remediation and validation of Langflow deployments are strongly recommended.

Technical Details

- Six critical vulnerabilities have been disclosed affecting Langflow OSS versions 1.0.0 through 1.10.0.
- CVE-2026-10134 is an unauthenticated remote code execution vulnerability with a CVSS score of 10.0, allowing attackers to execute arbitrary code without authentication.
- CVE-2026-7803 is a flow validation bypass vulnerability with a CVSS score of 9.8, potentially enabling unauthorized execution or manipulation of application workflows.
- CVE-2026-7871 is an insecure Redis deserialization vulnerability with a CVSS score of 9.8 that may facilitate arbitrary code execution through malicious serialized data.
- CVE-2026-7873 is a code injection and operating-system command execution vulnerability with a CVSS score of 9.9, allowing execution of attacker-controlled commands on affected systems.
- CVE-2026-10140 is a cross-tenant API credential reuse vulnerability with a CVSS score of 9.6, potentially enabling access to resources belonging to other tenants.
- CVE-2026-7663 is a Streamable MCP authorization bypass vulnerability with a CVSS score of 9.8 according to NVD and 9.1 according to IBM CNA.
- Collectively, the vulnerabilities may allow arbitrary code execution, authorization bypass, application integrity compromise, sensitive data access, and complete takeover of vulnerable Langflow environments.
- The flaws affect AI and LLM application development environments, potentially exposing AI workflows, integrated services, credentials, and connected data sources.
- Proof-of-concept exploits have been demonstrated for several of the vulnerabilities, increasing the likelihood of exploitation attempts.
- No confirmed active exploitation has been reported at the time of disclosure; however, the severity and availability of exploit material elevate the overall risk profile.

Recommendations

- Immediately upgrade all Langflow OSS deployments to version 1.10.1 or later.

Vulnerability and affected product details can be found *here* and *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
SimpleHelp Authentication Bypass Vulnerability Actively Exploited	**HIGH**	**CLEAR**	**Vulnerability**	**Open Source**

Executive Summary

SimpleHelp has disclosed a critical authentication bypass vulnerability, CVE-2026-48558, affecting SimpleHelp Remote Monitoring and Management (RMM) servers configured to use OpenID Connect (OIDC)

authentication. The flaw allows unauthenticated attackers to submit forged identity tokens and bypass authentication controls, potentially obtaining technician-level access to vulnerable environments without valid credentials. The vulnerability carries a CVSS 3.1 score of 10.0 and affects both legacy and pre-release product versions. Organizations utilizing internet facing SimpleHelp deployments face elevated risk due to the potential for complete unauthorized access to managed systems.

Organizations in the financial sector should be aware that active exploitation has been observed in the wild. Threat actors have reportedly leveraged the vulnerability to deploy the TaskWeaver Node[.]js malware loader and Djinn Stealer, enabling credential theft, cloud-access compromise, browser data collection, cryptocurrency-wallet theft, and theft of AI-development tokens. These capabilities could affect organizations by facilitating unauthorized access, downstream compromise of managed endpoints, and broader intrusion activity across enterprise environments.

Technical Details

- CVE-2026-48558 is a critical authentication bypass vulnerability affecting SimpleHelp deployments configured to use OpenID Connect (OIDC) authentication.
- The vulnerability allows unauthenticated attackers to bypass authentication by submitting forged identity tokens.
- Successful exploitation can result in unauthorized technician-level access to vulnerable SimpleHelp servers.
- All 6.0 pre-release versions configured with OIDC authentication are also affected.
- Active exploitation has been observed targeting vulnerable SimpleHelp environments.
- Threat actors have been observed deploying the TaskWeaver Node[.]js malware loader following successful compromise.
- Subsequent activity includes deployment of Djinn Stealer, which is capable of harvesting credentials and sensitive information from compromised systems.
- Information targeted by the malware includes account credentials, cloud-access keys, browser data, cryptocurrency-wallet information, and AI-development tokens.

Recommendations

- Immediately upgrade vulnerable SimpleHelp deployments to version 5.5.16, 6.0 RC2, or later supported releases.
- Treat exposed and unpatched SimpleHelp servers as potentially compromised until validated otherwise.
- Rotate credentials, API keys, cloud-access credentials, administrative passwords, and authentication tokens associated with affected environments.

Vulnerability and affected product details can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Adobe Campaign Classic Security Update Addresses Remote Code Execution Risk	HIGH	CLEAR	Vulnerability	Open Source

Executive Summary

Adobe has released security updates for Adobe Campaign Classic to address a critical vulnerability that could allow arbitrary code execution on affected systems. The vulnerability, tracked as CVE-2026-48286, stems from an incorrect authorization flaw and carries a CVSS v3.1 score of 10.0, indicating maximum severity. The issue affects on-premises Adobe Campaign Classic deployments and could permit complete compromise of vulnerable instances if successfully exploited.

Organizations in the financial sector should be aware that Adobe Campaign Classic is often integrated with customer engagement, marketing, and communication workflows containing sensitive customer and business data. Successful exploitation could affect the confidentiality, integrity, and availability of affected environments. While Adobe has stated that it is not aware of active exploitation in the wild, the critical severity warrants immediate remediation.

Technical Details

- Adobe published security bulletin APSB26-69 on 30 June 2026 addressing a critical vulnerability in Adobe Campaign Classic.
- The flaw is classified as Incorrect Authorization (CWE-863).
- Successful exploitation could result in arbitrary code execution on affected systems.
- The vulnerability affects on-premises Adobe Campaign deployments, including fully on-premises environments and on-premises components used within hybrid deployments.
- Adobe-hosted Campaign Classic environments have already been remediated and require no customer action.
- Adobe has assigned the update a Priority 1 rating, indicating a high-priority security update requiring prompt deployment.
- Adobe stated that it is not aware of any exploits in the wild targeting this vulnerability at the time of publication.

Recommendations

- Immediately upgrade affected Adobe Campaign Classic installations to version 7.4.3 build 9397 or later.

Vulnerability and affected product details can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Apache Tomcat Authentication Bypass Vulnerability	MEDIUM	CLEAR	Vulnerability	Open Source

Executive Summary

The Apache Software Foundation has disclosed CVE-2026-55957, an authentication bypass vulnerability affecting multiple supported versions of Apache Tomcat. The vulnerability occurs when security constraints are configured for the default servlet, causing configured HTTP method restrictions or method omissions to be ignored under affected conditions. This behavior may allow unauthorized access to protected resources and weaken intended access-control enforcement.

Organizations in the financial sector should be aware that web applications relying on Tomcat security constraints for resource protection could be exposed to authentication and authorization bypass risks. Exploitation may affect the integrity of access-control mechanisms protecting sensitive business applications and services.

Technical Details

- The vulnerability is classified by Apache as an Important security issue.
- The issue affects deployments where security constraints are specified for the default servlet.
- In vulnerable versions, any configured HTTP method restrictions or method omissions may be ignored, resulting in unintended access-control behavior.
- The vulnerability may enable unauthorized access to resources that administrators intended to protect using servlet security constraints.

Recommendations

- Immediately upgrade affected Apache Tomcat installations to the latest supported fixed release.

Vulnerability and affected product details can be found *here*.

back to top

Appendix A - Tactics, Techniques & Procedures (TTPs)

Fake AI-Agent Threats Drive Infostealer Delivery, Brand Impersonation, and Autonomous Intrusion Activity

Tactic	Technique	Name	Campaign Usage
Resource Development	T1583.001	Acquire Infrastructure: Domains	Registration of AI-vendor look-alike domains (macos-claude[.]com, ai[.]deepseekem[.]com, setup-code[.]com, payments[.]openai-services[.]org)
Resource Development	T1587.001	Develop Capabilities: Malware	Fake model repositories and trojanized AI installers built within hours of a product launch; susano vulnerable-driver loader
Initial Access	T1189	Drive-by Compromise	Malvertising and SEO-poisoned AI-tool download pages and fake AI video-generator sites (UNC6032)
Execution	T1204.002	User Execution: Malicious File	Victim runs a trojanized “Claude Setup.exe” / fake DeepSeek installer
Execution	T1203	Exploitation for Client Execution	AI-orchestrated autonomous exploitation of target applications (GTG-1002 agent + tool servers)
Credential Access	T1555.003	Credentials from Web Browsers	Infostealer harvest of browser-stored credentials, cookies, and wallet data
Command and Control	T1102	Web Service	EtherHiding blockchain-hosted C2 and messaging-platform exfiltration; Chaos C2 (deepseeklab[.]xyz, 45[.]153[.]186[.]237)
Resource Development	T1588.001	Obtain Capabilities: Malware	Commodity infostealers (Amatera, AMOS, Vidar, Rhadamanthys, Lumma) and the Chaos botnet bundled into AI-tool lures
Initial Access	T1566.002	Phishing: Spearphishing Link	AI-vendor payment and login phishing (payments[.]openai-services[.]org) and AI-branded adversary-in-the-middle credential capture
Execution	T1059.001	Command and Scripting Interpreter: PowerShell	ClickFix lures induce victims to paste attacker PowerShell/terminal commands piping a remote script to a shell
Defense Evasion	T1036.005	Masquerading: Match Legitimate Name or Location	Binaries and sites impersonate AI-vendor brand names and visual identity; CIS keyboard-layout abort checks
Credential Access	T1539	Steal Web Session Cookie	Session-token theft via AMOS/Amatera and AI-branded adversary-in-the-middle phishing
Exfiltration	T1567.002	Exfiltration to Cloud Storage/Web Service	Stealer-log exfiltration to underground messaging channels; UAE credentials surface in combo-list and stealer-log markets

ToddyCat Leverages OAuth Authorization Workflow for Email Access

Tactics	Techniques	Observed Activity
Defense Evasion	T1574.001 DLL Side-Loading	Umbrij was executed through DLL sideloading using legitimate signed applications vulnerable to insecure DLL loading.
Privilege Escalation / Defense Evasion	T1134.003 Access Token Manipulation: Make and Impersonate Token	The malware duplicated tokens from user processes to obtain and operate within the target user context.
Discovery	T1082 System Information Discovery	The tool searched the system for browser installations, user profiles, and local application data repositories.

Discovery	T1217 Browser Information Discovery	Umbrij enumerated browser profiles and identified authenticated accounts associated with Google services.
Collection	T1213 Data from Information Repositories	The malware copied browser databases, storage repositories, profile settings, synchronization information, and credential-related data.
Credential Access	T1555 Credentials from Password Stores	The tool collected browser files containing saved account and credential-related information.
Initial Access / Credential Access	T1528 Steal Application Access Token	The operation was designed to acquire OAuth authorization codes and exchange them for access tokens to access Google resources.
Collection	T1119 Automated Collection	Umbrij automated the collection and processing steps required for account access and token acquisition.

ARToken Expands Device Code Phishing Operations Against Microsoft 365

Tactics	Techniques	Observed Activity
Initial Access	Phishing	ARToken used targeted vendor-impersonation and invoice-themed phishing emails to initiate the compromise chain.
Credential Access	Steal Application Access Token	The platform abused Microsoft 365 device authorization workflows to capture authentication tokens and bypass MFA.
Persistence	Account Manipulation	ARToken supported Primary Refresh Token persistence to maintain access after initial compromise.
Collection	Email Collection	The operator panel exposed capabilities for accessing compromised Microsoft 365 mailboxes.
Collection	Data from Information Repositories	ARToken included functionality to access and collect data from SharePoint repositories.
Exfiltration	Exfiltration Over Web Service	The platform supported SharePoint data exfiltration as part of its post-compromise toolkit.
Defense Evasion	Obfuscated Files or Information	The phishing kit used encrypted payload delivery as part of its anti-analysis system.
Defense Evasion	Virtualization or Sandbox Evasion	ARToken used client-side behavioral verification to detect or avoid analysis environments.
Impact	Financial Theft	The platform supported BEC operations, including invoice-themed lures and mailbox abuse that could enable fraudulent payment activity

Appendix B – Threat Severity Ratings & Definitions

In this newsletter, each campaign, breach, and vulnerability has been assigned a severity level—High, Medium, or Low—based on a contextual assessment of its impact, relevance to the financial sector, regional exposure, and exploit status.

A **High** severity rating is assigned to threats that involve widely used technologies, nation-state actors, or actively exploited zero-day vulnerabilities. These incidents pose a significant risk due to their potential for widespread disruption, data compromise, and strategic targeting of core infrastructure.

A **Medium** severity rating reflects threats that are technically impactful but have limited regional exposure or affect less prominent entities. These may include ransomware attacks on mid-sized firms, blockchain-based malware delivery mechanisms, or smart contract exploits in DeFi platforms.

A **Low** severity rating is applied to threats that are peripherally related to the financial sector, have no confirmed impact in the region, or target non-critical systems. These include credential stuffing attacks on platforms outside the MENA region or supply chain breaches that have not yet affected local infrastructure.

Threat Score Ratings & Definitions

1. **Severe**: Widespread, high-impact attacks ongoing. Critical infrastructure affected by destructive malware or multiple zero-day exploits.
2. **High**: Confirmed breaches or coordinated campaigns in progress. High-impact threat actor activity and zero-day exploitation observed.
3. **Elevated**: Active exploitation of known vulnerabilities or targeted campaigns detected. Some breaches or zero-days may surface.
4. **Guarded**: Slight increase in malicious activity. Low-level campaigns or threat actor reconnaissance may be underway.
5. **Normal**: No significant cyber activity beyond baseline. No major breaches, campaigns, or zero-day activity observed.

Appendix C – Traffic Light Protocol (TLP) Definitions and Usage

TLP	When should it be used?	How should it be shared?
TLP:Red	Sources may use TLP:RED when information cannot be effectively acted upon without significant risk for the privacy, reputation, or operations of the organizations involved. For the eyes and ears of individual recipients only, no further.	Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed. In the context of a meeting, for example, TLP:RED information is limited to those present at the meeting. In most circumstances, TLP:RED should be exchanged verbally or in person.
TLP:Amber+Strict	Sources may use TLP:AMBER+STRICT when information requires support to be effectively acted upon, yet carries risk to privacy, reputation, or operations if shared outside of the organization.	Recipients may share TLP:AMBER+STRICT information only with members of their own organization on a need-to-know basis to protect their organization and prevent further harm.

TLP:Amber	Sources may use TLP:AMBER when information requires support to be effectively acted upon, yet carries risk to privacy, reputation, or operations if shared outside of the organizations involved. Note that TLP:AMBER+STRICT should be used to restrict sharing to the recipient organization only.	Recipients may share TLP:AMBER information with members of their own organization and its clients on a need-to-know basis to protect their organization and its clients and prevent further harm.
TLP:Green	Sources may use TLP:GREEN when information is useful to increase awareness within their wider community.	Recipients may share TLP:GREEN information with peers and partner organizations within their community, but not via publicly accessible channels. Unless otherwise specified, TLP:GREEN information may not be shared outside of the cybersecurity or cyber defense community.
TLP:Clear	Sources may use TLP:CLEAR when information carries minimal or no foreseeable risk of misuse, in accordance with applicable rules and procedures for public release.	Recipients may share this information without restriction. Information is subject to standard copyright rules.

Appendix D - Acronyms & Technical Terms

Term / Acronym	Meaning / Description
AAA	Authentication, Authorization and Accounting. A framework used for identity verification and access management.
Agentic AI	AI systems that can independently perform multi-step actions using connected tools and workflows.
AI	Artificial Intelligence. Technology that enables systems to perform tasks that typically require human intelligence.
AI Agent	An autonomous or semi-autonomous AI system capable of performing actions, interacting with tools, and executing tasks on behalf of users.
AI Supply Chain	The ecosystem of AI models, agents, plugins, skills, tool servers, and dependencies used by AI-enabled environments.
AI-Generated Content	Content created by artificial intelligence models, including code, text, images, or lures used in cyber operations.
AI-Orchestrated Intrusion	An attack in which AI agents perform portions of the intrusion lifecycle with limited human involvement.
AiTM	Adversary-in-the-Middle. A phishing technique used to intercept credentials and session tokens during authentication.
Amatera	An information-stealing malware family frequently delivered through AI-themed lures.
AMOS	Atomic macOS Stealer. A malware family designed to steal credentials, cookies, session data, and cryptocurrency wallet information from macOS systems.

Anubis	A ransomware-as-a-service operation whose affiliates were observed exploiting remote access infrastructure, stealing credentials, and deploying ransomware.
API	Application Programming Interface. A mechanism that allows software applications to communicate and exchange data.
Application Access Token	A token used to authorize access to applications or cloud resources without directly using passwords.
APSB	Adobe Product Security Bulletin. Adobe's advisory format for security vulnerabilities and patches.
Armored Likho	A threat group targeting government and critical infrastructure organizations using phishing and information-stealing malware.
ARToken	A phishing-as-a-service platform used to conduct Microsoft 365 device-code phishing and token theft operations.
AsyncRAT	A remote access trojan that enables attackers to control infected systems and execute malicious actions remotely.
BEC	Business Email Compromise. Fraudulent activity involving compromised email accounts to conduct financial or operational deception.
BITS	Background Intelligent Transfer Service. A Windows component commonly abused by malware for file transfers.
Browser Extension	Software add-on that extends browser functionality and may have access to browsing activity and data.
Browser Hijacking	Modification of browser settings or traffic redirection without legitimate user intent.
BusySnake	A Python-based information stealer used in the Armored Likho campaign.
BusySnake Stealer	A Python-based information stealer used by Armored Likho to collect sensitive data from compromised systems.
C2 / Command-and-Control	Infrastructure used by attackers to manage compromised systems and receive stolen data.
CanisterWorm	A self-propagating worm used to spread malicious code through software packages and development environments.
CAPTCHA	Challenge-response mechanism intended to distinguish humans from automated systems.
CastleLoader	Malware loader used to deploy additional malicious payloads.
CI/CD	Continuous Integration / Continuous Delivery or Deployment. Automated software build and release processes.
CipherForce	A ransomware operation previously associated with TeamPCP before its partnership with Vect.
CitrixBleed 2	A NetScaler vulnerability enabling exposure of session materials and facilitating session hijacking.
ClickFix	A social-engineering technique that tricks users into manually executing malicious commands.
Cloud Credential	Authentication information used to access cloud services and resources.
Cloudflare Pages	A web-hosting service that was abused to distribute malicious verification pages and malware delivery infrastructure.
Cloudflare Tunnel / cloudflared	Software used to create secure tunnels between internal systems and external services.
Combo List	A collection of usernames and passwords typically used for credential stuffing and account compromise.

ConnectWise ScreenConnect	A legitimate remote support tool abused by threat actors for persistence and remote access.
Credential Harvesting	Collection of usernames, passwords, tokens, and authentication material from victims or systems.
Credential Stuffing	Automated abuse of stolen username-password pairs across multiple services.
Cross-Tenant Attack	An attack where resources, credentials, or data belonging to one tenant are improperly accessed by another tenant.
CSC	UAE Cyber Security Council
CVE	Common Vulnerabilities and Exposures. A publicly tracked identifier assigned to security vulnerabilities.
CVPN	Clientless Virtual Private Network functionality available in gateway deployments.
CVSS	Common Vulnerability Scoring System. A framework used to measure vulnerability severity.
CWE	Common Weakness Enumeration. A standardized catalog of software security weaknesses.
Data Leak Site	Website operated by ransomware groups to publish stolen data and pressure victims into paying extortion demands.
DeBank	A decentralized finance portfolio management platform whose brand was impersonated in an AI-focused typosquatting campaign.
DeclarativeNetRequest (DNR)	Chromium browser functionality that allows extensions to intercept, redirect, or filter web requests.
Device Code Phishing	An attack abusing OAuth device authentication workflows to obtain access tokens from users.
DevSecOps	Development, Security, and Operations practices integrated into software delivery processes.
Digital Surveillance	Monitoring and collection of digital information about users, systems, or organizations.
Djinn Stealer	Malware used to steal credentials, browser data, cryptocurrency wallet information, cloud keys, and other sensitive information.
DLL	Dynamic Link Library. A shared Windows software component that can be abused through DLL sideloading techniques.
DLL Sideloading	A technique where a legitimate executable loads a malicious library placed in an expected location.
DNS	Domain Name System. Infrastructure used to translate domain names into network addresses.
DNS Proxy	A service that forwards DNS requests on behalf of clients.
DoS	Denial of Service. An attack that disrupts system availability or performance.
EDR	Endpoint Detection and Response. Security technology used to detect and respond to threats on endpoints.
EtherHiding	A technique using blockchain infrastructure to host or retrieve malicious configurations.
EvilTokens	A phishing-as-a-service ecosystem focused on Microsoft 365 device-code phishing and token theft.
Exfiltration	Unauthorized extraction or transfer of data from a compromised system.
Ferocious Kitten	A previously reported surveillance group whose malware tradecraft overlaps with activity associated with MarkiRAT.

FIPS	Federal Information Processing Standards. Security standards used in regulated technology products.
Flow Validation Bypass	A weakness allowing security checks governing workflows or application logic to be bypassed.
Franz	A legitimate messaging application that was trojanized in one malware delivery chain.
GitHub Repository	A code storage location frequently used for software development and sometimes abused for malware hosting.
GSSAPI	Generic Security Services Application Program Interface. A framework supporting secure authentication mechanisms.
GTG-1002	A reported AI-orchestrated cyber-espionage campaign leveraging AI agents and connected tools.
GTIG	Google Threat Intelligence Group. Google's threat intelligence and cyber threat tracking organization.
Hallucinated Domain	A fictitious web domain generated by an AI model that could later be weaponized by attackers.
HijackLoader	Malware loader used to deploy additional payloads such as information stealers and remote access malware.
HTTP/2	A modern web communications protocol used by web servers and applications.
ICA Proxy	Citrix feature used to broker virtual application and desktop sessions.
Identity Token	Authentication artifact used to verify a user's identity to applications and services.
IdP	Identity Provider. A service that authenticates users and provides identity assertions to applications.
Indirect Prompt Injection (IPI)	Hidden instructions embedded within content designed to manipulate AI-agent behavior.
Infostealer	Malware designed to collect credentials, tokens, cookies, and other sensitive information.
Inline Content Analysis	Security inspection of web content or files as they are processed or delivered.
JNDIRealm	Apache Tomcat component used for authentication against directory services.
JSON-LD	JavaScript Object Notation for Linked Data. Structured metadata frequently used by websites and search engines.
JWT	JSON Web Token. A token format commonly used for authentication and authorization.
Langflow	Open-source platform used to build AI and LLM-powered applications.
Lateral Movement	Attack activity used to move from one compromised system to others within an environment.
LiteLLM	Widely used AI gateway software that was compromised during a supply-chain attack.
LLM	Large Language Model. An AI model trained to generate or process natural language.
LNK File	Windows shortcut file that can be weaponized to execute malicious commands.
Lumma	Information-stealing malware distributed through various malware campaigns.
Malvertising	Use of malicious advertisements to lure victims to phishing pages or malware.

MCP	Model Context Protocol. A framework used to connect AI systems with tools and services.
MeshAgent	Legitimate remote monitoring and management software abused by attackers for persistence.
MFA	Multi-Factor Authentication. Authentication requiring multiple verification factors.
Mimikatz	A credential-extraction tool frequently used by attackers to access authentication material.
Montana Empire	Phishing kit referenced in the phantom-squatting research demonstrating weaponization of AI-hallucinated domains.
NetScaler ADC	Citrix Application Delivery Controller appliance used for load balancing, remote access, and application delivery.
NetScaler Gateway	Citrix remote-access solution used for secure application and VPN access.
NSIP	NetScaler management IP address used for administration.
OAuth 2.0	Authorization framework that enables delegated access to applications and services.
OAuth Token	Authentication token that grants authorized access to cloud services without repeatedly entering credentials.
OIDC	OpenID Connect. An authentication protocol built on OAuth 2.0.
Oracle Load Balancer	Network traffic-distribution deployment mode affected by one of the NetScaler vulnerabilities.
Out-of-Bounds Read	A vulnerability allowing software to access memory outside intended boundaries.
Password Export	Extraction of passwords stored within browsers or applications.
Perplexity AI	AI-powered search platform whose branding was impersonated by a malicious browser extension.
PhaaS	Phishing-as-a-Service. Criminal platforms that provide phishing infrastructure to affiliates.
Phantom Domain	A non-existent domain name generated by an AI model that may later be registered and weaponized by attackers.
Phantom Squatting	Registration and weaponization of AI-hallucinated domains to exploit trust in AI-generated recommendations.
PoC	Proof of Concept. Demonstration code or exploit showing vulnerability exploitability.
Process Hollowing	Malware technique where malicious code executes within a legitimate process.
Prompt Injection	An attack designed to manipulate the behavior of AI models by embedding malicious instructions in content processed by the model.
PRT	Primary Refresh Token. Authentication token used to maintain access to cloud resources.
PsExec	Administrative tool often abused for remote execution and lateral movement.
PyArmor	Python code-obfuscation technology used to hinder malware analysis.
RaaS	Ransomware-as-a-Service. Criminal model where ransomware operators provide tooling to affiliates.
RAG	Retrieval-Augmented Generation. AI process combining model responses with retrieved external information.
RAT	Remote Access Trojan. Malware that enables remote control of infected devices.

RCE	Remote Code Execution. The ability to execute arbitrary code on a target system.
RDP	Remote Desktop Protocol. Technology for remote system access commonly targeted by attackers.
React2Shell	Critical software vulnerability previously exploited by TeamPCP to compromise development infrastructure.
Redis	In-memory data store often used by applications and AI platforms.
Remote Access Tool (RAT)	Software that enables remote control of systems, either for legitimate administration or malicious purposes.
Remotely	Legitimate remote administration software observed being abused by ransomware operators.
Remus	Information-stealing malware distributed in ClickFix campaigns.
Repository Poisoning	Malicious modification of software repositories to distribute compromised code or packages.
ResiLoader	Malware loader observed delivering additional payloads including infostealers.
RMM	Remote Monitoring and Management. Administrative software frequently abused by threat actors.
SAML	Security Assertion Markup Language. Protocol used for identity federation and authentication.
Secure Private Access (SPA)	Hybrid remote-access deployment architecture that may incorporate NetScaler infrastructure.
SEO Poisoning	Manipulation of search-engine results to increase visibility of malicious websites.
Session Token	Authentication artifact used to maintain access to services after login.
SimpleHelp	Remote Monitoring and Management platform affected by a critical authentication bypass vulnerability.
Slopsquatting	Registration of AI-hallucinated package names or dependencies for supply-chain attacks.
SMB	Server Message Block. Network file-sharing protocol often used during lateral movement.
Social Engineering	Psychological manipulation techniques used to persuade users to perform actions that benefit attackers.
SOCKS Proxy	A network proxy service used to relay communications through another system.
SSO	Single Sign-On. Authentication process allowing access to multiple services using one login.
StealC	Information-stealing malware commonly delivered through malware distribution campaigns.
STRD	Shadow Token via Remote Debug. OAuth token acquisition technique leveraging browser remote debugging.
Streamable MCP	Langflow component affected by an authorization bypass vulnerability.
Supply Chain Attack	Compromise of software, services, dependencies, or providers to gain access to downstream victims.
Supply Chain Compromise	Breach of a trusted software, service, or provider that impacts downstream users and organizations.
TaskWeaver	Node[.]js malware loader observed following exploitation of SimpleHelp servers.
TCP Timestamp	Networking feature used in TCP communications and referenced in a NetScaler vulnerability.

TDS	Traffic Distribution System. Infrastructure used to direct victims to malicious content.
TeamPCP	Threat group known for supply-chain compromises, credential theft, and collaboration with ransomware operators.
Tenant Isolation	Security separation between customers or environments in a shared platform.
THP	Thermal Hallucination Persistence. Measurement of consistency in AI-generated hallucinated domains.
Threat Actor	An individual, group, or organization conducting malicious cyber activity.
ToddyCat	Advanced threat group observed conducting email-access and OAuth-token theft operations.
Token	A digital authentication or authorization artifact used to access systems or services.
Token Theft	Theft of authentication tokens that can be used to access accounts and services without passwords.
Trivy	Open-source vulnerability scanner that was compromised during a documented software supply-chain attack.
Trojanized Installer	Legitimate-looking software package modified to deliver malware.
TTPs	Tactics, Techniques, and Procedures. The methods and behaviors used by threat actors during attacks.
Typosquatting	Registration of look-alike domains designed to impersonate legitimate brands or services.
UAC	User Account Control. Windows security mechanism that limits unauthorized administrative actions.
Umbrij	Custom tool used by ToddyCat to obtain OAuth authorization tokens through browser remote debugging functionality.
Vect	Ransomware-as-a-service operation that partnered with TeamPCP to monetize stolen credentials through ransomware deployment.
Verification Lure	Fraudulent prompt or message designed to convince users to perform an action such as executing code or providing credentials.
VPN	Virtual Private Network. Technology enabling secure remote access to networks.
VPS	Virtual Private Server. Cloud-hosted server frequently used as attack infrastructure.
Workflow Manipulation	Unauthorized modification of automated business, application, or AI-driven processes
Worm	Malware capable of self-propagating across systems and networks without direct user action.
XOR	Exclusive OR. A basic encryption or obfuscation method commonly used by malware.
YEPlayer	Example of a fake application used in MarkiRAT malware delivery operations.
Zero-Click Prompt Injection	Prompt-injection technique requiring no direct interaction from the user to influence an AI agent.
Zero-Reputation Domain	Newly registered domain with little or no intelligence history, making detection more difficult for security controls.
Zoho Assist	Legitimate remote access software observed being abused for persistence and remote control.