

Anomali Cyber Watch: F5 Breach, Mysterious Elephant APT, Malicious MCP Servers, and More

By Anomali Cyber Watch

Published: 2025-10-21 · Archived: 2026-07-22 02:00:36 UTC


All Posts

11

min read

Published on

- [Breach of F5 Requires Emergency Action From BIG-IP Users](#)
- [TA585 Self Managed Campaigns Leverage “ClickFix” and MonsterV2](#)
- [Mysterious Elephant APT Evolves Its Arsenal](#)
- [Jewelbug APT Executes Five-Month Espionage Against Russian IT Provider](#)
- [When Trusted AI Connections Turn Hostile](#)
- [LinkPro Linux Rootkit Exposes Deep Stealth Threat to Cloud & Linux Environments](#)
- [Silver Fox Expands Winos 4.0 Attacks to Japan and Malaysia via HoldingHands RAT](#)
- [LastPass Warns of Phishing Emails Impersonating Company Breach Notice](#)
- [Fake Homebrew/LogMeIn Sites Push Info Stealers](#)
- [Pixnapping Attack Exposes Android On-Screen Secrets](#)



[Breach of F5 Requires Emergency Action From BIG-IP Users](#)

(published: October 16, 2025)

F5 disclosed that a sophisticated nation-state actor gained long-term access to its internal development and engineering systems, stealing portions of BIG-IP source code, undisclosed vulnerability data, and a subset of customer configuration files. In response, the US Cybersecurity and Infrastructure Security Agency (CISA) issued Emergency Directive 26-01, ordering federal agencies to inventory, patch, or disable affected F5 products and to restrict public exposure of management interfaces. F5 reported that independent audits found no signs of source-code tampering or supply-chain insertion, and no known exploitation of undisclosed vulnerabilities so far.

However, over 200,000 exposed F5 systems have been identified online globally, underscoring the urgency for rapid remediation and network isolation.

Analyst Comment: The real concern here is that the attackers didn't just steal data, they took insight into how F5 products work and where the weak spots might be. That kind of access means they could quietly develop exploits long after the headlines fade. Even though there's no sign of active abuse yet, it's smart to assume they're preparing for it. Anyone running BIG-IP should treat this as a forward-looking risk, not a past incident. Patch quickly, lock down management interfaces, and monitor for anything unusual.

MITRE ATT&CK: [T1078 - Valid Accounts](#) | [T1005 - Data From Local System](#) | [T1041 - Exfiltration Over C2 Channel](#) | [T1070 - Indicator Removal On Host](#)

TA585 Self Managed Campaigns Leverage “ClickFix” and MonsterV2

(published: October 14, 2025)

Cybersecurity researchers have identified a sophisticated crime group known as TA585, which manages its entire attack chain from phishing delivery to malware installation without relying on third-party infrastructure. The group distributes the MonsterV2 malware family, a combined stealer, loader, and remote access trojan, using phishing lures that impersonate U.S. government entities and through JavaScript injections on compromised websites. These sites display fake CAPTCHA prompts, a technique called “ClickFix,” that tricks users into running PowerShell commands which deliver the payload. TA585 has also been observed abusing GitHub notifications to distribute malicious links. MonsterV2 is capable of credential theft, screenshot and webcam capture, hidden remote access, cryptocurrency clipping, and loading additional payloads. It is frequently wrapped with the SonicCrypt crypter to evade detection. Command-and-control activity has been seen on port 7712 across multiple samples in recent campaigns.

Analyst Comment: By running every part of their operation themselves, they remove weak links that usually expose or slow down threat actors. That control lets them adapt fast, test new delivery methods, and keep campaigns running even when individual domains or payloads get taken down. The ClickFix trick shows the same mindset, simple, consistent, and built to scale rather than impress. Defenders should pay attention to patterns in infrastructure reuse and staging behavior rather than chasing single payloads, because that's where TA585's advantage really lies.

MITRE ATT&CK: [T1566.002 - Phishing: Spearphishing Link](#) | [T1189 - Drive-By Compromise](#) | [T1204.001 - User Execution: Malicious Link](#) | [T1059.001 - Command and Scripting Interpreter: Powershell](#) | [T1027 - Obfuscated Files Or Information](#) | [T1497.001 - Virtualization/Sandbox Evasion: System Checks](#) | [T1622 - Debugger Evasion](#) | [T1555.003 - Credentials from Password Stores: Credentials From Web Browsers](#) | [T1113 - Screen Capture](#) | [T1125 - Video Capture](#) | [T1005 - Data From Local System](#) | [T1095 - Non-Application Layer Protocol](#) | [T1573 - Encrypted Channel](#) | [T1041 - Exfiltration Over C2 Channel](#)

Mysterious Elephant APT Evolves Its Arsenal

(published: October 17, 2025)

A newly documented campaign in early 2025 shows that the Mysterious Elephant advanced persistent threat (APT) group has shifted from reusing known malware to deploying customized and heavily modified tools. The group primarily targets government and foreign affairs organizations in the Asia-Pacific region, especially Pakistan, Bangladesh, Sri Lanka, Nepal, and Afghanistan. Initial access is typically achieved through spear-phishing emails carrying weaponized Office documents exploiting CVE-2017-11882. Post-compromise activity involves custom PowerShell loaders, the BabShell C++ reverse shell, and the MemLoader HidenDesk reflective loader. The group also exfiltrates data from WhatsApp and uses stealthy infrastructure tactics like wildcard DNS, per-request subdomains, and cloud-hosted servers. Indicators of compromise include file hashes, domains, and IPs tied to the campaign.

Analyst Comment: Mysterious Elephant's shift toward developing its own loaders and reflective tools shows a clear level-up in capability. They're no longer leaning on recycled malware but refining tradecraft quietly to stay under the radar. The use of WhatsApp data exfiltration is a calculated move, taking advantage of trusted platforms to disguise espionage activity. This level of evolution suggests the group may be ready to move beyond regional operations, potentially testing new geographies or targeting higher-value sectors as its confidence and resources grow.

MITRE ATT&CK: [T1203 - Exploitation For Client Execution](#) | [T1059.001 - Command and Scripting Interpreter: Powershell](#) | [T1620 - Reflective Code Loading](#) | [T1560 - Archive Collected Data](#) | [T1041 - Exfiltration Over C2 Channel](#)

Jewelbug APT Executes Five-Month Espionage Against Russian IT Provider

(published: October 16, 2025)

China-linked advanced persistent threat (APT) group Jewelbug, also tracked as CL-STA-0049, REF7707, and Earth Alux, carried out a five-month espionage campaign against a Russian IT service provider between January and May 2025. The operation is notable as a rare case of cross-nation targeting between strategic partners. Jewelbug used a renamed Microsoft debugger (cdb.exe disguised as 7zup.exe) to execute shellcode, bypass allowlisting, and disable endpoint defenses. The group performed credential dumping, established persistence through scheduled tasks, and wiped Windows Event Logs to cover its tracks. Data exfiltration was conducted via Yandex Cloud using a custom tool, yandex2.exe, blending with normal network traffic. Access to software build systems and code repositories suggests possible supply-chain compromise. Jewelbug has been active since at least 2023, favoring long dwell times and stealth over speed.

Analyst Comment: Jewelbug's targeting of a Russian IT provider is both unusual and strategically telling. Chinese and Russian threat actors have historically avoided each other's operational domains, making this a rare breach of that unspoken boundary. The introduction of a custom backdoor in this campaign shows deliberate investment in tailored tooling built for persistence, stealth, and data staging within the victim's environment. Defenders should monitor for renamed system binaries, unfamiliar executables, and subtle deviations in normal process behavior. Jewelbug's focus on blending with trusted infrastructure means that static signatures alone won't suffice, behavioral monitoring and strong endpoint visibility are key to detection.

MITRE ATT&CK: [T1218 - Signed Binary Proxy Execution](#) | [T1053.005 - Scheduled Task/Job: Scheduled Task](#) | [T1003 - Os Credential Dumping](#) | [T1070.001 - Indicator Removal on Host: Clear Windows Event Logs](#) | [T1567.002 - Exfiltration Over Web Service: Exfiltration To Cloud Storage](#)

When Trusted AI Connections Turn Hostile

(published: October 16, 2025)

Malicious Model Context Protocol (MCP) servers present a significant and under-recognized threat to AI systems. These servers, which facilitate connections between large language models (LLMs) and external tools, data, or services, can be weaponized to seize host control, manipulate outputs, and exfiltrate sensitive information, often with no detectable signature. Researchers demonstrated twelve classes of attacks targeting configuration, metadata, prompt logic, tool behavior, and more, with many achieving full success. Traditional MCP scanners detected only a handful of malicious servers in tests. This issue is compounded by open deployment practices: thousands of MCP servers are publicly published without vetting. The risk is no longer theoretical, security firms discovered a malicious “postmark-mcp” package that silently forwarded emails.

Analyst Comment: MCP security will only improve when it is treated as part of the broader supply chain rather than a side component of LLM integration. Borrowing from OWASP frameworks is a practical start, but real progress comes from making MCP oversight a routine part of development and deployment. Every link in that chain, from the code to the connectors and the data they handle, needs verification and visibility. The researchers are right that responsibility is shared. Auditing, validation, and runtime checks must become standard practice, not optional safeguards. The sooner organizations apply the same scrutiny to MCP servers as they do to production APIs, the fewer surprises they will face as AI systems become more deeply embedded in critical operations.

MITRE ATT&CK: [T1048 - Exfiltration Over Alternative Protocol](#) | [T1020 - Automated Exfiltration](#) | [T1082 - System Information Discovery](#) | [T1552 - Unsecured Credentials](#) | [T1078 - Valid Accounts](#) | [T1059 - Command And Scripting Interpreter](#)

Silver Fox Expands Winos 4.0 Attacks to Japan and Malaysia via HoldingHands RAT

(published: October 18, 2025)

The China-linked threat actor known as Silver Fox has broadened its operations beyond China and Taiwan, now targeting Japan and Malaysia with its remote-access Trojan campaign. The group is deploying its established malware family, Winos 4.0 (also called ValleyRAT), alongside a new payload dubbed HoldingHands RAT (aka Gh0stBins). The attack chain typically begins with phishing emails that carry PDF attachments masquerading as official finance or tax documents, these either embed links or redirect victims to malicious download sites. Once the initial payload executes, it uses techniques like DLL sideloading, anti-virtual-machine checks, privilege escalation, and termination of security tools. The malware sets up persistence via scheduled tasks, drops files to C:\Windows\System32 and disables the Windows Task Scheduler service’s recovery mechanism. The infrastructure includes cloud buckets and command-and-control servers hosting encrypted modules disguised as

image files. Earlier campaigns by Silver Fox targeted Taiwan and used trojanised medical-viewer software and pivoted via SEO poisoning and gaming-application decoys.

Analyst Comment: Silver Fox's move into Japan and Malaysia looks deliberate, not experimental. They are sticking with their dependable Winos 4.0 framework while introducing the new HoldingHands RAT, which shows a team that evolves its tools rather than replacing them. The continued use of tax-themed lures and legitimately signed binaries points to access to stolen or privately held certificates, something that deserves closer attention. What stands out to me is their reliance on Alibaba OSS for hosting payloads. Using a regional cloud service gives their traffic a layer of authenticity that blends into local network patterns. It feels less like a campaign expansion and more like a refinement of an established playbook designed to last.

MITRE ATT&CK: [T1566.001 - Phishing: Spearphishing Attachment](#) | [T1059.001 - Command and Scripting Interpreter: Powershell](#) | [T1053 - Scheduled Task/Job](#) | [T1068 - Exploitation For Privilege Escalation](#) | [T1574.002 - Hijack Execution Flow: Dll Side-Loading](#) | [T1218.005 - Signed Binary Proxy Execution: Mshta](#) | [T1105 - Ingress Tool Transfer](#) | [T1082 - System Information Discovery](#) | [T1113 - Screen Capture](#) | [T1041 - Exfiltration Over C2 Channel](#)

LastPass Warns of Phishing Emails Impersonating Company Breach Notice

(published: October 13, 2025)

LastPass issued an alert about a phishing campaign attempting to trick users into believing the company had been hacked. The emails, sent from "hello@lastpasspulse[.]blog" and "hello@lastpassgazette[.]blog," use the subject line "We Have Been Hacked - Update Your LastPass Desktop App to Maintain Vault Security." Recipients are directed to fake domains including "lastpassdesktop[.]com," "lastpassgazette[.]blog," and "lastpassdesktop[.]app," all hosted via bulletproof provider NICENIC. The sites prompt users to download a fake desktop application that steals credentials. The campaign was launched over the US holiday weekend, likely to delay detection. LastPass confirmed its systems remain secure and is working with Cloudflare and hosting providers to disable the fraudulent infrastructure.

Analyst Comment: The phishing emails appear convincing, using credible sender domains and urgent wording to exploit users' lingering concern over a potential LastPass breach. The actor registered several look-alike domains and hosted them through NICENIC, timing the campaign for a US holiday weekend to delay detection. It's a calculated, low-effort operation built on social engineering rather than technical sophistication. The key point is awareness: LastPass is not distributing any desktop app updates by email, and any link making that claim is part of this ongoing phishing campaign.

MITRE ATT&CK: [T1566.002 - Phishing: Spearphishing Link](#) | [T1204.001 - User Execution: Malicious Link](#) | [T1555 - Credentials From Password Stores](#) | [T1583.001 - Acquire Infrastructure: Domains](#) | [T1027 - Obfuscated Files Or Information](#)

Fake Homebrew/LogMeIn Sites Push Info Stealers

(published: October 18, 2025)

A new campaign is targeting macOS users, especially developers, by advertising fraudulent versions of trusted platforms such as Homebrew, LogMeIn, and TradingView via sponsored links on major search engines. These ads lead to cloned download portals instructing users to paste a Terminal command. That command downloads and executes one of two advanced info-stealers: AMOS (Atomic macOS Stealer) or Odyssey Stealer. The malware evades macOS protection mechanisms such as Gatekeeper, gains elevated privileges, harvests credentials, browser cookies, crypto-wallets, and other sensitive data, then exfiltrates it to attacker-controlled infrastructure. Researchers found more than 85 malicious domains tied to the campaign, many linked by shared certificates and infrastructure.

Analyst Comment: AMOS has been appearing in more campaigns recently, and this one reinforces that trend. The technique of advertising fraudulent versions of trusted platforms is also becoming more common, especially through paid search results. It's an effective way to reach users who already trust the names being impersonated and are likely to follow install instructions without question. This feels less like opportunism and more like a deliberate refinement of social engineering for macOS users.

MITRE ATT&CK: [T1566 - Phishing](#) | [T1059 - Command And Scripting Interpreter](#) | [T1555.001 - Credentials from Password Stores: Keychain](#) | [T1041 - Exfiltration Over C2 Channel](#)

Pixnapping Attack Exposes Android On-Screen Secrets

(published: October 14, 2025)

A newly disclosed side-channel vulnerability in Android devices, called Pixnapping, allows malicious applications to secretly capture on-screen information from other apps such as two-factor authentication codes, private messages, and location data. The attack takes advantage of GPU timing behavior, known as GPU.zip, together with Android's rendering APIs, which lets unprivileged apps infer screen content pixel by pixel without any special permissions. It has been demonstrated on devices including Google Pixel 6 through 9 and Samsung Galaxy S25 running Android 13 to 16. The issue is tracked as CVE-2025-48561. Google released partial mitigations in the September 2025 Android Security Bulletin, but researchers confirmed that workarounds remain possible and a complete fix is still pending.

Analyst Comment: Pixnapping is a proof of concept exploit that exposes a major gap in mobile security. It works without any special permissions and shows that even what appears on a phone's screen can be stolen, including 2FA codes that people rely on to protect their accounts. This shifts attention from data stored or sent over networks to the information shown to users in real time. It also underlines why mobile devices deserve the same level of protection and monitoring as traditional endpoints. While this is not an active attack in the wild, it offers a glimpse of how future threats may target the visual layer itself rather than just apps or data.

MITRE ATT&CK: [T1113 - Screen Capture](#) | [T1041 - Exfiltration Over C2 Channel](#) | [T1202 - Indirect Command Execution](#) | [T1068 - Exploitation For Privilege Escalation](#)



December 24, 2025

Anomali Cyber Watch

Anomali Cyber Watch: Sample 3

LockBit 5.0 Ransomware Targets Windows, Linux, and VMware ESXi in Active Campaigns. Google Patches Actively Exploited Chrome Zero-Day CVE-2026-2441. Infostealer Targets OpenClaw Configuration Files to Capture Credentials and User Context. And more...

[Read More](#)



December 24, 2025

Anomali Cyber Watch

Anomali Cyber Watch: Sample

LockBit 5.0 Ransomware Targets Windows, Linux, and VMware ESXi in Active Campaigns. Google Patches Actively Exploited Chrome Zero-Day CVE-2026-2441. Infostealer Targets OpenClaw Configuration Files to Capture Credentials and User Context. And more...

[Read More](#)



December 24, 2025

Anomali Cyber Watch

Anomali Cyber Watch: Sample 2

LockBit 5.0 Ransomware Targets Windows, Linux, and VMware ESXi in Active Campaigns. Google Patches Actively Exploited Chrome Zero-Day CVE-2026-2441. Infostealer Targets OpenClaw Configuration Files to Capture Credentials and User Context. And more...

[Read More](#)

[Explore All](#)

Source: <https://www.anomali.com/blog/anomali-cyber-watch-f5-breach-mysterious-elephant-apt-malicious-mcp-servers>