


Suspected Chinese Operators Use Claude Code and DeepSeek to Breach Government Systems Across Four Countries

 hunt.io/blog/chinese-operators-claude-deepseek-government-intrusion

July 14, 2026



Disclosure note: The affected organizations and the relevant national CERTs were notified on July 6, 2026. Publication was held for a 7-day disclosure window.

In June 2026, a pivot from known TencShell C2 infrastructure led us to an open directory exposing an active intrusion campaign. TencShell is a Go-based implant derived from the open-source Rshell framework, first documented by [Cato CTRL in May 2026](#) and assessed there as suspected China-linked. The directory held victim source code, custom exploit scripts, operational logs, and cloned login pages, with the attack notes written in Simplified Chinese.

What caught our attention was the tooling behind it. Claude Code and DeepSeek-v4-pro ran as working parts of the intrusion, not tools off to the side. They handled reasoning for bypass techniques, reworked exploits after failed attempts, and built the phishing pages used to harvest credentials. That puts this campaign alongside [Anthropic's November 2025 disclosure](#) of a China-linked operation that used Claude Code to automate large-scale intrusions.

The following summarizes the findings from our research.

Key Findings

- The open directory on 112.213.124[.]132 exposed a full operational toolkit, including payloads, operator files and scripts, and victim-specific folders.
- Infrastructure pivots identified 13 Hong Kong-based servers spanning four separate ASNs, three with shared SSH keys and TLS certificates, serving in a redundancy capacity.
- Active exploitation of government systems in Afghanistan, Thailand, and Taiwan, with reconnaissance and phishing staging against U.S. government portals.
- Scanning of 5,890+ government hosts across 10 countries with a Python-developed automated scoring scale.
- Claude Code handled execution and session persistence while DeepSeek-v4-pro drove the reasoning, a two-model split documented across the recovered log files.

Our pivot began with a single network fingerprint.

Discovery and Infrastructure

Building on public threat research documenting [TencShell](#) C2 infrastructure, we pivoted on a HTTP header hash observed on port 1111 of the command and control nodes. Using HuntSQL, we developed a simple query to search for additional servers sharing identical header fingerprints over the last 10 days.

HuntSQL Query:

```
SELECT
  ip,
  port,
  min(timestamp),
  max(timestamp)
FROM
  httpv2
WHERE
  http.headers.redacted.hash.sha256 == "03f26cbfa3ca15fcb43f512aa4041732beeec267f9d1dc74a11f7b0bb32e86bb"
AND
  timestamp > NOW - 10 DAY
GROUP BY
  ip,
  port
```

Copy

The query returned 13 unique IP addresses, the remaining 11 servers, previously unreported, are all located in Hong Kong and distributed across four autonomous systems: VMISS Inc., MEGA-II IDC, CTG Server Limited, and Antbox Networks Limited. Beyond port 1111, the cluster also exposed services on ports 1212, and 8090.




Figure 01: HuntSQL [results](#) for the TencShell malware port 1111 HTTP header, showing 13 unique IP's (query date: 10 Jun 2026).

Among the IP's, [112.213.124.\[132\]](#) (MEGA-II IDC, AS152194) was observed hosting an open directory on port 8888. Additionally, the server was flagged as high risk for exposing Asset Reconnaissance Lighthouse ([ARL](#)) and a Vshell C2 service.




Figure 02: Hunt.io IP profile for 112.213.124.[.]132 showing open ports 1111, 3000 (DeepAudit), 5003 (ARL), 8084 (Vshell), and 8888 (open directory).

The following ports and services were running on the server at the time of analysis:

Port	Service	Purpose
22	SSH	Remote Access
1111	HTTP	Malware Download
3000	DeepAudit	Code Vulnerability Scanning
5003	ARL	Asset Reconnaissance
8084	Vshell	C2 Server
8888	HTTP	Open Directory

Table 1: Ports and services observed on 112.213.124.[.]132

The ports and services on this IP supported a complete attack workflow. ARL performed network reconnaissance, [DeepAudit](#), which we found running on port 3000, audits code for vulnerabilities, and Vshell provides command and control. ARL, Vshell, and DeepAudit are publicly available and have legitimate testing uses, so their presence alone does not confirm malicious activity. Here it is the open directory sitting alongside them, holding victim source code, exploit scripts, and operator logs, that ties the tooling to an active intrusion rather than authorized testing.

Infrastructure Pivoting

The Associations tab in Figure 02 revealed two additional servers tied to the open directory at 112.213.124.[.]132. Hosted on the same network block, 112.213.124.[.]159 and 112.213.124.[.]163 have shared an SSH host key fingerprint (64107E3E0A333F685D1BE6386426223A030C4126AC7C295AA7B1D54C508BBACE) with .132 since May 17, 2026. All three servers ran ARL on port 5003, serving an identical TLS certificate.

The certificate is the project's default that ships with every install. The pertinent fields are provided below:

- C = CN
- ST = Shanghai
- L = Shanghai

- CN = 127.0.0.1
- O = Example Inc.
- OU = Web Security
- SHA-256: AD1A0B3E22A10A2BD680B773B178A0D3824CFCBDF3551016F3D052A0B823079F

Between June 18-19, 2026, all three servers reissued their certificate, producing a new one that retained the defaults, but produced a fresh key and serial (SHA-256: 2954639BE599F23C2229A9743ABA09A1D9D11BF2BECC62BF353384437DB37DEE). This update points to a recent redeployment/restart of ARL by the operators, indicating maintenance and upkeep of the servers continues into mid-late June 2026.

IP	Company	Country	SSH Fingerprint	First Seen	Last Seen
112.213.124.159	CTG Server Limited	HK	dbf8d8dbc63a16afe0f9d21e26b62e6f6a03b25c3a98c16e9876edc043e514ec	2026-05-13	2026-06-12
112.213.124.163	CTG Server Limited	HK	dbf8d8dbc63a16afe0f9d21e26b62e6f6a03b25c3a98c16e9876edc043e514ec	2026-05-13	2026-06-12

Figure 03: Hunt.io Associations for [112.213.124.\[.\]132](#) showing two additional servers sharing SSH keys and TLS certificates. All three servers exposed the same services: an open directory on port 8888, a Vshell listener, ARL on port 5003, and DeepAudit. We reviewed the open directory contents on .132 and .159 directly; the matching services on .163 were identified by fingerprint.

HTTP GET requests to port 1111 on 112.213.124.[.]132 triggered the download of a previously unreported Linux/ARM 32-bit binary, HSEWH-Ur. The Golang-compiled, statically linked executable beacons over WebSocket to port 4081 on the same host, and uses X-NITRO-USER, X-NITRO-PASS for HTTP authentication headers. These headers are used by legitimate infrastructure-management APIs.

The other two servers delivered similar ARM variants on the same port, each communicating with 112.213.124.[.]132:4081, indicating this server acts as some sort of C2 hub.

Additional Capabilities:

- Tencent messaging platform (QQ/IM) credentials extraction, including SDK identifiers and cryptographic keys
- Exfil of enterprise messaging platform login details (corpId/corpSecret patterns)
- Cloud service access keys theft
- File upload/download

HSEWH-Ur and its counterparts are Linux/ARM samples recovered from the same infrastructure we pivoted through from the TencShell C2 nodes. We stop short of calling them Linux/ARM builds of TencShell itself. The implant Cato documented was a Windows sample that persisted through a Registry Run key, and we have not confirmed a code-level match between it and

these Linux files. The focus on consumer messaging, enterprise platforms, and cloud credentials fits malware built for cloud and edge device environments.

A similar sample was recovered from 38.55.105[.]143:8088, with some noticeable differences. The Linux/x86 build differs from the previously mentioned ARM files, and uses garble to strip function and variable names removing references to the Recon project seen previously. Despite these differences, both malware variants share an identical 80-byte encryption key, indicating at the very least a shared codebase across the two architectures. Static analysis pointed to the above IP:port combo as a possible C2, however no network communications were observed in sandbox analysis.

A Possible Second C2 Framework

Two of the thirteen servers, 192.229.115[.]229, and 192.229.115[.]230, served more than TencShell infrastructure. Both presented a TLS certificate on port 443 and 8083 self-identifying as a separate C2 framework. The subject and issuer fields consist of CN = Gshell Server and O = Gshell C2. We were unable to find prior public reporting documenting a framework under this name. Port 8089 hosts a generic Chinese-language login panel titled 'X'.

Pivoting on the certificate subject fields surfaced 7 servers (after deduplicating ports) not including the two mentioned above.

HuntSQL Query:

```
SELECT
  ip,
  port
FROM
  certificates
WHERE
  subject.common_name = 'Gshell Server'
  AND subject.organization = 'Gshell C2'
GROUP BY
  ip,
  port
```

Copy

The results from our query showed similar location and hosting providers we saw with TencShell. All servers were located in Hong Kong, across Antbox Networks Limited and CTG Server limited.

The screenshot displays the Hunt.io SQL Explorer interface. On the left is a navigation sidebar with sections for Platform, Hunting, and Integrations. The main area is titled 'SQL Editor' and contains a SQL query. Below the editor, the 'Results' section shows a table with 14 rows of data. The table has two columns: 'ip' and 'port'. The results are grouped by IP address, with each IP having two associated ports listed.

ip	port
192.163.16730	8083
192.163.16730	443
192.229.115.230	8083
192.229.115.230	443

Figure 04: HuntSQL [output](#) querying the Gshell TLS certificate fields revealing 9 total unique IPs.

Because .229 and .230 belong to both the TencShell and Gshell clusters, the two frameworks appear to be operated in tandem, rather than independently. Although we were unable to recover malware samples, we assess with moderate confidence that Gshell represents a second C2 used by the operators.

Beyond the hosted tooling and malware, the open directories left a far more valuable footprint exposed; operator logs, scripts, and more.

Directory Contents

Port 8888 on 112.213.124[.]132, running Python SimpleHTTP contained 2,431 files and 80 subdirectories, revealing multiple files and directories: JSP and PHP web shells (some disguised as images), custom exploitation scripts, cloned webpages and login portals, network scan output, and dumps of victim data. Logs documenting attack execution and coordination were also observed, and will be covered in a separate section.




Figure 05: Snippet of the attacker-controlled directory at 112.213.124[.]132:8888. Files targeting specific entities did not follow a consistent naming convention. Some used explicit organization names, others employed acronyms, and the remaining, shortened designators. This variance suggests the directory aggregates output from multiple operators working independently rather than an automated pipeline with standardized naming patterns.

Analyzing the files and folders within the directory provided visibility into the actors' targeting strategy and operational priorities.

That analysis broke down into a clear set of victims, starting with the staged phishing pages.

Targeting Scope & Victimology

Taiwan

Further review of the files and subfolders revealed targeting across Taiwanese supply chain and manufacturing sectors. Reconnaissance activity spanned 8 organizations, with varying levels of probing. From what we were able to identify, two organizations in critical industries were successfully exploited, with timestamps indicating activity as recent as 11 June 2026.

The organizations below were enumerated and fingerprinted only. We saw no exploitation against any of them. We have described them by sector rather than by name, since they were reconnaissance targets and not confirmed compromises:

Sector	Recon activity
Container shipping operator	DNS and subdomain enumeration, password brute-forcing
Container shipping operator	Domain enumeration

This asset mapping provided a foundation for the activity that followed against both the chemical manufacturer and telecommunications company.

Thailand

A Thai government administrative system fell victim to a SQL injection exploitation using SQLMap. The attackers achieved authentication bypass and admin panel access, enabling deployment of a GIF polyglot webshell for persistent command execution via GET parameters.

ที่	เลขประจำตัว	ชื่อ	ตำแหน่ง	เงินเดือน	กลุ่ม	รูปภาพ	สถานะ	แก้ไข
1							สำเร็จ	แก้ไข
2							สำเร็จ	แก้ไข
3	666666666666	test shell test	ลูกจ้าง				สำเร็จ	แก้ไข
4	111111111111	test test php5	ลูกจ้าง				สำเร็จ	แก้ไข
5	999999999999	test hacker test	ลูกจ้าง				สำเร็จ	แก้ไข
6	555555555555	test shellpng test	ลูกจ้าง				สำเร็จ	แก้ไข
7	111111111111	test test pht	ลูกจ้าง				สำเร็จ	แก้ไข
8	777777777777	test hackdotphp test	ลูกจ้าง				สำเร็จ	แก้ไข
9	111111111111	test test inc	ลูกจ้าง				สำเร็จ	แก้ไข
10	888888888888	test hacknull test	ลูกจ้าง				สำเร็จ	แก้ไข
11	111111111111	test test shtml	ลูกจ้าง				สำเร็จ	แก้ไข
12							สำเร็จ	แก้ไข
13							สำเร็จ	แก้ไข
14							สำเร็จ	แก้ไข
15							สำเร็จ	แก้ไข
16							สำเร็จ	แก้ไข
17							สำเร็จ	แก้ไข
18							สำเร็จ	แก้ไข
19							สำเร็จ	แก้ไข
20							สำเร็จ	แก้ไข
21							สำเร็จ	แก้ไข

Figure 07: Screenshot of the redacted compromised database table of a Thai government admin panel. The exfiltrated database contained the names, National ID numbers and positions of government employees. Test entries created by the attackers during the intrusion indicated hands-on access and active data manipulation.

Admin panel access confirmed activity as recent as June 9, 2026. The directory contained 980 files referencing this system alone, indicating both a lengthy compromise and a focus on this particular government office.

From the files we recovered, the above represented the only activity targeting Thailand.

Afghanistan

An Afghan government application was compromised, exposing both citizen complaint submissions and core application infrastructure. Exfiltrated source code appeared in two directories: [victim]_git/ and [victim]_git_dump/, revealing a Laravel framework version 5.8.38 installation. Encryption keys, database credentials, and mail handling code were also dumped. Using the recovered credentials, operators developed a custom Python exploit targeting Laravel's deserialization mechanisms to achieve remote code execution.

```

1 #!/usr/bin/env python3
2 """
3 Laravel APP_KEY RCE Exploit for gov.af
4 Target: Laravel 5.8.38, Monolog 1.27.0
5 APP_KEY: base64:
6 """
7
8 import base64
9 import json
10 import hmac
11 import hashlib
12 import struct
13 import os
14 import sys
15 import urllib.parse
16 from Crypto.Cipher import AES
17 import requests
18 import warnings
19 warnings.filterwarnings('ignore')
20 from urllib3.exceptions import InsecureRequestWarning
21 requests.packages.urllib3.disable_warnings(InsecureRequestWarning)
22
23 # ===== CONFIG =====
24 TARGET = "https://gov.af"
25 APP_KEY_B64 = "
26 COMMAND = "id" # Default command
27
28 # ===== LARAVEL ENCRYPTER =====
29 class LaravelEncrypter:
30     def __init__(self, key_b64):
31         self.key = base64.b64decode(key_b64)
32         self.cipher = 'AES-256-CBC'
33     def encrypt(self, plaintext):

```

Figure 08: Snippet of the exploit code targeting an Afghan government web application.

Additional activity observed included: CSRF tokens extracted from live user sessions, and file upload parameter enumeration. Six distinct versions of the complaint submission form were copied to the directory, showing a specific focus on this feature of the web app.

The database credentials provided access to citizen-submitted complaints and grievances. Combined with the recovered mail infrastructure code and live application authentication, attackers would have the ability to monitor reporting channels in real-time and identify individuals and institutions referenced in complaint data. For state actors, this represents high-value visibility into emerging unrest, and vulnerabilities across governmental organizations.

United States

The United States briefly appeared in the operation not through confirmed intrusions, as seen above, but at earlier stages of the attack lifecycle through reconnaissance and prepped phishing pages. At the network scanning level, logs recorded NASA hosts `launchpad.nasa.gov` and `ngis.nasa.gov`. No follow-on activity was identified against either, suggesting the hosts were enumerated during the broader scanning activity but not pursued further.

We were able to recover cloned web pages impersonating two U.S. government entities: the D.C. Council and the government of Delaware County, Pennsylvania. The pages were at differing levels of completion. The homepage for D.C. Council reproduced styling and formatting but was missing its images, while a matching clone of the site's WordPress login page was fully built.

The Delaware County clone, which specifically copied the county web pages contact form, was also unfinished, missing visual elements from the legitimate site.





Figure 09: Cloned admin login page spoofing the D.C. Council website.

The disparity between a polished login page and a half-rendered homepage suggests the operators prioritized the credential-capturing pages, and development was likely still in progress. No client-side credential stealing code was found, leaving the possibilities that this is handled server-side, or the pages had not yet reached that stage in being operationalized.

As mid-tier government administrative bodies, D.C. Council and Delaware County handle procurement and contracting activity, vendor relationships, local policies, and constituent correspondence. The targeting of a county contact form is consistent with interest in the latter, a means that provides citizens, businesses, and organizations the ability to raise grievances with the local government. This selection aligns with documented Chinese intelligence collection priorities: mapping contractor and supply chain systems and gaining early insight into government policy and procurement.

Financial Services Targeting

In addition to the government-sector activity, the operators ran a parallel campaign against financial services firms across multiple regions. The clearest example being an attacker-developed CORS exploit page on 112.213.124[.]159 that successfully extracted WordPress administrator account data from a large payment processing platform. A cross-reference on the exposed accounts against public LinkedIn profiles, confirmed individuals with the same name as employees of the company.



The supporting files and folders: enumeration scripts, logs and screenshots documenting failed password brute-force attempts showed a specific targeting of billing platforms across Europe, Australia, and Asia. The combination of this material alongside the government targeting on the same dedicated infrastructure implies the financial services attacks as a second set of objectives for compromise.

Attacker Methodology: LLM Integration

Claude Code version 2.1.165 conducted operations across multiple persistent sessions. Session IDs show the same infrastructure was used across different geographic targeting campaigns, with Taiwan-specific operations saved to dedicated working directories. The timestamps span from June 8 - 12, 2026.

Beyond attack planning, Claude Code managed the attackers' phishing infrastructure operations. Instructions in a recovered CLAUDE.md file directed it to create, test, and iterate cloned pages, automatically refining the pages across multiple targets.

Threat Actor Assessment

11/13

Indicators of Compromise

File Hashes

Filename	Algorithm	Hash	Data
HSEWH-Ur	SHA-256	90b7b2c6f3d05234dc55678243039d7e51f0d54190239e5234a0005533337dc8	Retrieved from 112.213.124[.]132:1111
8eA-GlbK	SHA-256	90b7b2c6f3d05234dc55678243039d7e51f0d54190239e5234a0005533337dc8	Retrieved from 112.213.124[.]159:1111
r4l3DqLA	SHA-256	90b7b2c6f3d05234dc55678243039d7e51f0d54190239e5234a0005533337dc8	Retrieved from 112.213.124[.]163:1111
Ar70qICi	SHA-256	643de2a1cf9148b896efecf560c9476fa56118ec477c4e15eb5c2da4b318061f	Retrieved from 38.55.105[.]143:8088

Network Infrastructure - TencShell

Type	Indicator	Role	Hosting
IP:Port	112.213.124[.]163:1111	Malware Delivery/Open Directory	MEGA-II IDC, HK
IP:Port	192.238.134[.]166:1212	Suspected TencShell Infrastructure	Antbox Networks Limited, HK
IP:Port	134.122.200[.]153:443	Suspected TencShell Infrastructure	CTG Server Limited, HK
IP:Port	192.229.115[.]229:8090	Suspected TencShell Infrastructure/Gshell	Antbox Networks Limited, HK
IP:Port	45.64.52[.]245:1111	Suspected TencShell Infrastructure	CTG Server Limited, Hong Kong
IP:Port	134.122.200[.]154:443	Suspected TencShell Infrastructure	CTG Server Limited, HK
IP:Port	112.213.124[.]159:1111	Malware Delivery/Open Directory	MEGA-II IDC, HK
IP:Port	45.64.52[.]242:1111	Cato Networks IoC	CTG Server Limited, HK
IP:Port	38.55.105[.]143:8088	Malware Delivery	VMISS Inc., HK
IP:Port	45.64.52[.]246:1111	Suspected TencShell Infrastructure	CTG Server Limited, HK
IP:Port	112.213.124[.]132:1111	Malware Delivery/Open Directory	MEGA-II IDC, HK
IP:Port	134.122.200[.]155:443	Suspected TencShell Infrastructure	CTG Server Limited, HK
IP:Port	192.229.115[.]230:8090	Suspected TencShell Infrastructure/Gshell	Antbox Networks Limited, HK

Network Observables - Gshell

Type	Indicator	Role	Hosting
IP	192.163.167[.]5	Gshell Certificate Match	Antbox Networks Limited, HK

Type	Indicator	Role	Hosting
IP	192.163.167[.]7	Gshell Certificate Match	Antbox Networks Limited, HK
IP	134.122.200[.]114	Gshell Certificate Match	CTG Server Limited, HK
IP	134.122.200[.]115	Gshell Certificate Match	CTG Server Limited, HK
IP	192.163.167[.]6	Gshell Certificate Match	Antbox Networks Limited, HK
IP	192.163.167[.]10	Gshell Certificate Match	Antbox Networks Limited, HK
IP	134.122.200[.]116	Gshell Certificate Match	CTG Server Limited, HK

Summary

The pieces here point in one direction. Deliberate targeting of government systems in Afghanistan, Thailand, and Taiwan, hands-on exploitation rather than spray-and-pray, and a parallel run at financial services across Europe, Australia, and Asia. Simplified Chinese throughout the code and notes, Hong Kong infrastructure spread across four ASNs, and a clear interest in procurement and supply chain visibility all line up with China-based activity, though we stop short of naming a group.

The part worth sitting with is the role the models played. Claude Code drove execution while DeepSeek-v4-pro handled the reasoning, and the recovered files show both touching reconnaissance, exploit development, and phishing across separate geographic campaigns. This wasn't a tool bolted on at the end. It was wired into the workflow.

→ **The open directory that exposed all of this came from a single infrastructure pivot. If you want to see how we track this kind of attacker infrastructure before it goes dark, [book a free demo](#) and we'll walk you through the pivots behind this campaign.**