etugen.io

July 23, 2026

THREAT HUNTING REPORT

HUNT AUTONOMOUS INFRASTRUCTURE
BEFORE IT HUNTS YOU.

Contents

INTRODUCTION 3

Naming Mappings 5

AT A GLANCE 7

EXECUTIVE SUMMARY 10

PART 1: THREAT ACTOR PROFILE 11

AI Agent Architecture 12

Timeline 13

PART 2: CAMPAIGN ANALYSIS 14

TECHNICAL INFRASTRUCTURE - C2 17

TECHNICAL INFRASTRUCTURE - TOOLING 18

VICTIM ANALYSIS 19

MITRE ATT&CK MAPPING 20

INDICATORS (IOC) 21

DEFENSE RECOMMENDATIONS 22

ETG-2026-002 · CONCLUSION 23

Introduction

Artificial intelligence no longer merely generates content; it is now taking over the **execution layer** of the attack cycle. This report examines an eCrime actor that runs its entire operation, from reconnaissance through exploitation to persistence, via an **autonomous LLM agent**. The actor automates CVE monitoring, PoC harvesting, target scanning and exploit execution without human intervention, driven by commands issued from a messaging application.

This report is based on first-hand analysis of a working directory that the actor **left exposed to the internet through its own OPSEC failure** (an open directory); it is not a simulation or theoretical assessment. No active access was made to any target system; all findings were obtained through passive collection via the **etugen.io** platform.

The subject actor **SYNTHETIC PHANTOM** exposed a **97,973-file** operator workspace across two open ports on IP address 87.58.201.80. The contents reveal the autonomous AI agent framework, the exploit arsenal, target lists and, critically, the actor's own live SSH private key.

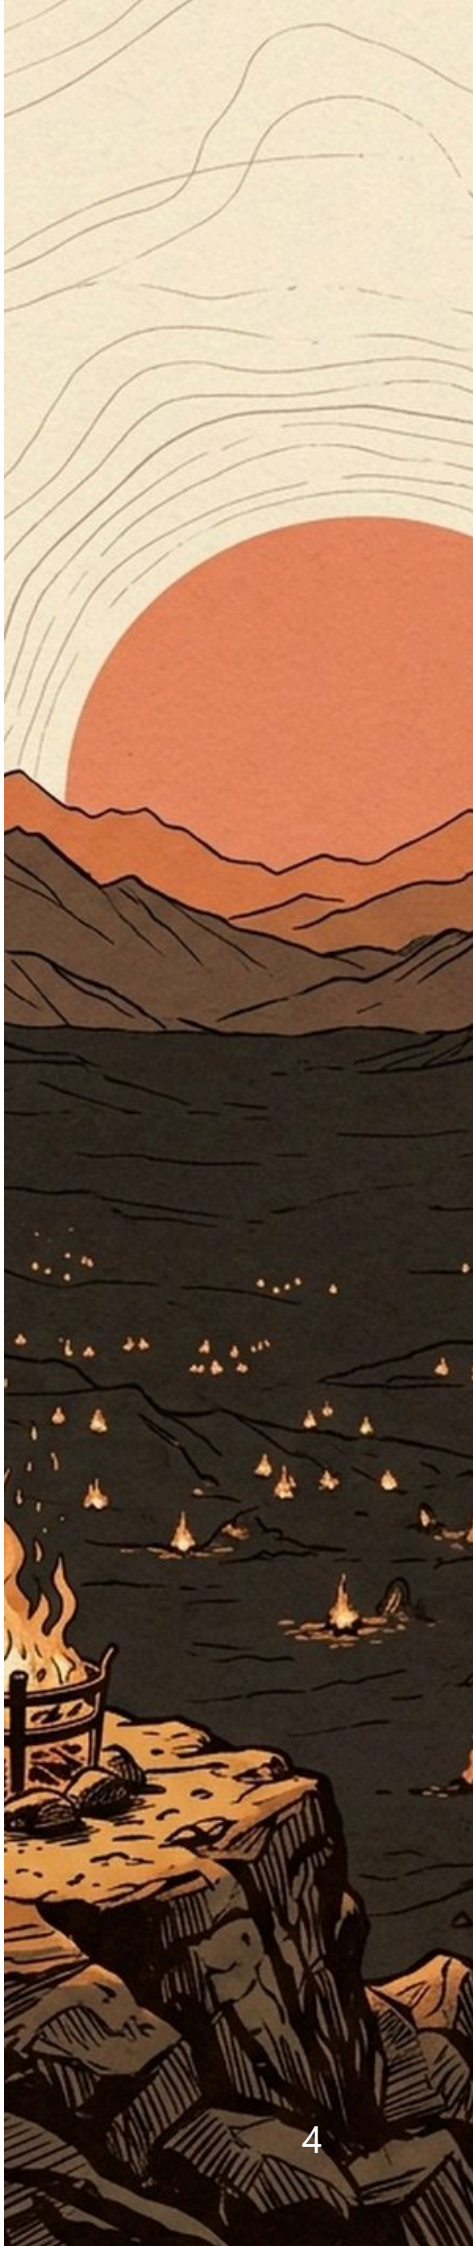
SYNTHETIC PHANTOM'S AUTONOMOUS AGENT AUTOMATED THE ATTACK CYCLE END TO END, FROM RECONNAISSANCE THROUGH POC HARVESTING TO MASS EXPLOITATION. THE ACTOR OBTAINED PERSISTENT ACCESS THROUGH A TARGETED ROOT COMPROMISE, A KYC-FRAUD REPOSITORY OF STOLEN FRENCH IDENTITY DOCUMENTS BELONGING TO DOZENS OF REAL INDIVIDUALS, AND THREE ORIGINAL 0-DAYS THAT BYPASS THE VENDOR PATCH.



AI-assisted attack automation presents a two-sided threat. On one hand, a single actor can concurrently scan and attempt to exploit hundreds of targets; on the other, this scale causes the actor to accumulate its own tooling, target lists and identity traces in a single place.

This report covers the operation actively observed between 2026-07-10 and 2026-07-15. During the analysis, the etugen.io Threat Intelligence Center documented the following findings:

- Live observation of the autonomous AI agent framework (~1,634 unique agent transcripts, skill/memory/kanban databases)
- Source-code-level detection of a custom module that automatically harvests PoCs from GitHub (User-Agent **HermesAutolearn/2.0**)
- Verification of the actor's three original 0-days (with Russian comments) that bypass the vendor patch
- Detection of an automated KYC-registration pipeline (register_and_exploit.py + OCR) that processes stolen genuine identity documents
- Detection of a persistent root backdoor obtained via container escape on the target server, reinforced by cron
- Exposure of the actor's live SSH private key and stolen French identity documents in the open directory



Naming Mappings

The **etugen.io** team assigns tracked threat actors a two-part code name: the adjective describes the actor's targeting focus or operational behavior; the category suffix indicates the actor's overall group. This system is designed to prevent the confusion caused by the same actor being referred to under different names across different sources.

THREAT ACTOR	CLASS	DESCRIPTION
PHANTOM	eCrime	Financially motivated cybercrime
SHADOW	APT	Nation-state affiliated crime
GHOST	HACKTIVIST	Ideological, politically motivated crime
VIPER	INSIDER THREAT	Abuse of trusted access

› ***etugen.io*** *code names do not imply the identity, nationality or country affiliation of actors. A code name is assigned to facilitate analysis and to avoid using the aliases the actor has chosen for itself.*

Actors in This Report

ALIAS	ATTRIBUTE	SUFFIX	DESCRIPTION
SYNTHETIC PHANTOM	SYNTHETIC	PHANTOM	Operational behavior of running the attack cycle via an autonomous AI agent (SYNTHETIC) + financially motivated eCrime category (PHANTOM)

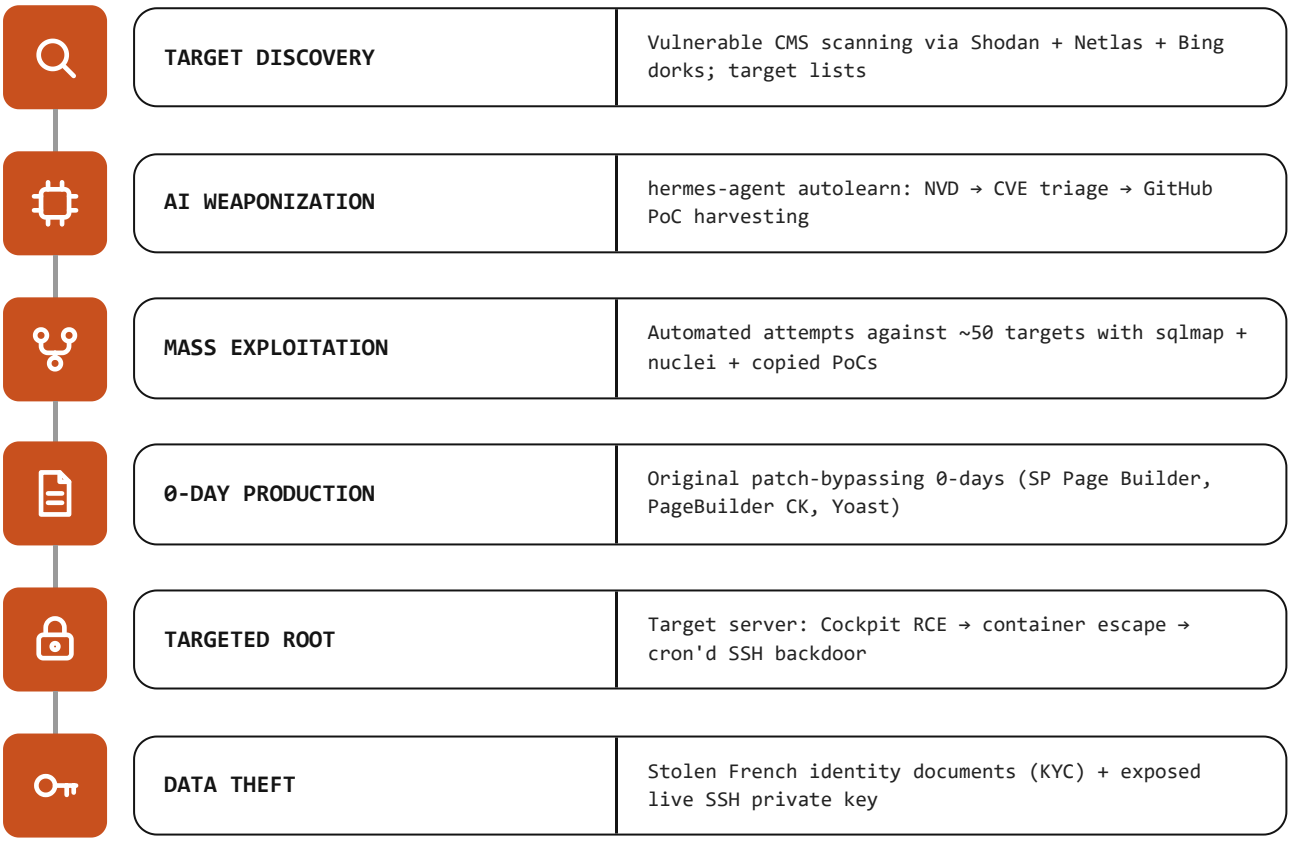


› ***etugen.io*** *code names do not imply the identity, nationality or country affiliation of actors. A code name is assigned to facilitate analysis and to avoid using the aliases the actor has chosen for itself (e.g. the agent name).*

At a Glance

This eCrime operation, executed by an AI agent, was observed through analysis of the open directory at 87.58.201.80. The findings document the scale a single operator reached through AI orchestration, and the concrete gains obtained at that scale: targeted root access and stolen identity data.

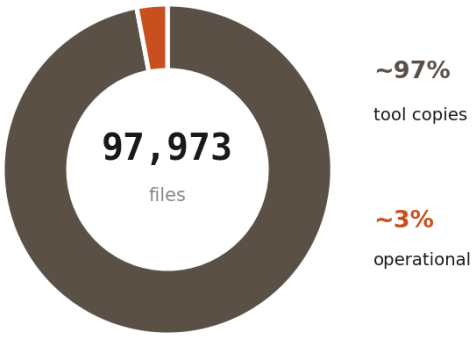
Anatomy of the Synthetic Phantom Attack



Open Directory Composition

The two open ports (7777 and 9876) are largely two copies of the same content. The overwhelming majority (~97%) of the **97,973 files** are **copies of publicly available source trees** such as ExploitDB, PoC-in-GitHub, nuclei-templates and WordPress/Asterisk; most of the hundreds of malware/APT labels the platform generates arise from this static noise. The real operational signal lies in the remaining ~3%.

This distinction is critical: the actor's **real capability** is visible not in thousands of copied exploits, but in three 0-days written with its own autonomous agent code and a single targeted root compromise. The analysis separates noise from operational evidence by file path.

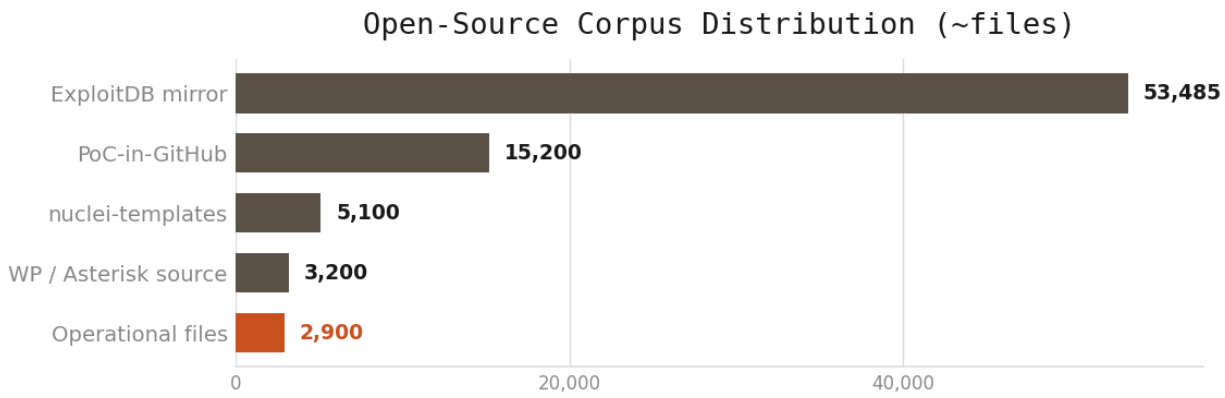


97,973
EXPOSED FILES (2 PORTS)

~1,634
UNIQUE AI AGENT TRANSCRIPTS

5,318
PYTHON EXPLOIT FILES

Corpus Distribution: Noise vs Signal



Access and Data Obtained

The actor's open directory documents the autonomous agent's **concrete gains** first-hand. The operation's real value lies not in noisy mass scanning; it concentrates in targeted persistent access, stolen identity data and original exploitation capability.

1
VERIFIED PERSISTENT ROOT BREACH (TARGET SERVER)

dozens
STOLEN GENUINE IDENTITY DOCUMENTS (FRENCH KYC)

3
ORIGINAL 0-DAYS BYPASSING THE VENDOR PATCH

OBTAINED	OPERATIONAL VALUE
Persistent root + cron'd SSH backdoor	Indefinite, self-healing admin access on the target server
Stolen French identity document repository	Directly monetizable PII for KYC fraud
Three original 0-days (Russian comments)	Patch-bypassing, reusable exploitation capability
Autonomous AI agent framework	High-tempo, scalable operation with a single operator

› *While the mass CMS scan remained broad but low-yield, the actor's real impact concentrates in **the persistent root it established through targeted manual work and the dozens of stolen identity documents**; in victim notification, priority belongs to these concrete gains.*

Executive Summary

This report examines the 2026-07-10 → 07-15 operations of a financially motivated eCrime actor we have code-named **SYNTHETIC PHANTOM**. The actor uses an Ubuntu workstation at 87.58.201.80 (QWINS LTD / AS213702) as its operations hub.

> **SYNTHETIC PHANTOM** runs the attack cycle via an autonomous LLM agent named "hermes-agent": it monitors CVEs from NVD, harvests PoCs from GitHub, and mass-scans with sqlmap/nuclei. The distinguishing trait is not sophistication but end-to-end **automation scale**.

Key Findings

- Autonomous AI agent (skill + persistent memory + kanban); operated via instant messaging by the operator
- GitHub PoC harvesting module (HermesAutolearn/2.0) collects repos by stars for a fixed CVE list
- Three actor-produced 0-days (SP Page Builder, PageBuilder CK, Yoast) read and bypass the vendor patch
- The autonomous agent reaches a tempo and scale that let a single operator scan hundreds of targets concurrently
- Verified persistent root on the target server via container escape + cron'd SSH backdoor
- The actor's live SSH private key and stolen French identity documents exposed in the open directory

Affected Areas: Global CMS sites (opportunistic), French individuals (identity/KYC), target server (root breach)

MITRE ATT&CK: T1595, T1588.006, T1190, T1059.006, T1203, T1505.003, T1611, T1053.003, T1552.001, T1567

Part 1: **Threat Actor Profile – SYNTHETIC PHANTOM**

SYNTHETIC PHANTOM is an eCrime actor operated by a single operator or a small team. It conducts its operations on a root-privileged Ubuntu server and uses an LLM agent as its primary execution layer.

Attribute	Assessment
Motivation	Financial (eCrime): identity fraud, crypto, mass webshell
Language (assessment)	Russian-speaking; Russian comments in the actor-produced 0-day code (medium-high confidence)
Targeting	French individuals (identity) + global CMS sites (opportunistic, mass scanning)
Technical Maturity	Medium: arsenal mostly copied PoCs; however original patch-bypass capability present
First Observed	2026-07-10 (open directory)
Last Activity	2026-07-15 (active at time of analysis)
Distinguishing Trait	LLM-orchestrated end-to-end autonomous attack; skill/memory/kanban architecture

The most notable feature of **SYNTHETIC PHANTOM** is that it runs the attack cycle through an autonomous AI agent. This agent:

- Runs on a large-language-model-based framework (persistent memory + kanban task tracking)
- Receives natural-language commands from the operator via instant messaging
- Independently performs CVE monitoring, PoC harvesting, target scanning and exploit execution
- Exhibits a human-directed, AI-executed model (verification_evidence.db present)

AI Agent Architecture

The operator manages the system with commands via an instant-messaging application. The LLM agent interprets the command, makes autonomous decisions and triggers the execution layers (reconnaissance, exploitation, exfil). This architecture lowers the operator's technical skill threshold and provides a much higher tempo than human operators.



`CVE Monitoring (NVD)` `PoC Harvesting (GitHub)` `Scanning & Exploitation`

Autolearn Pipeline

1

Phase 1 · `cve_monitor.py`

Monitors new CVEs from NVD (120 s loop).

2

Phase 2 · `content_finder.py`

Finds relevant PoCs and performs content triage.

3

Phase 3 · `collect_all_cves.py`

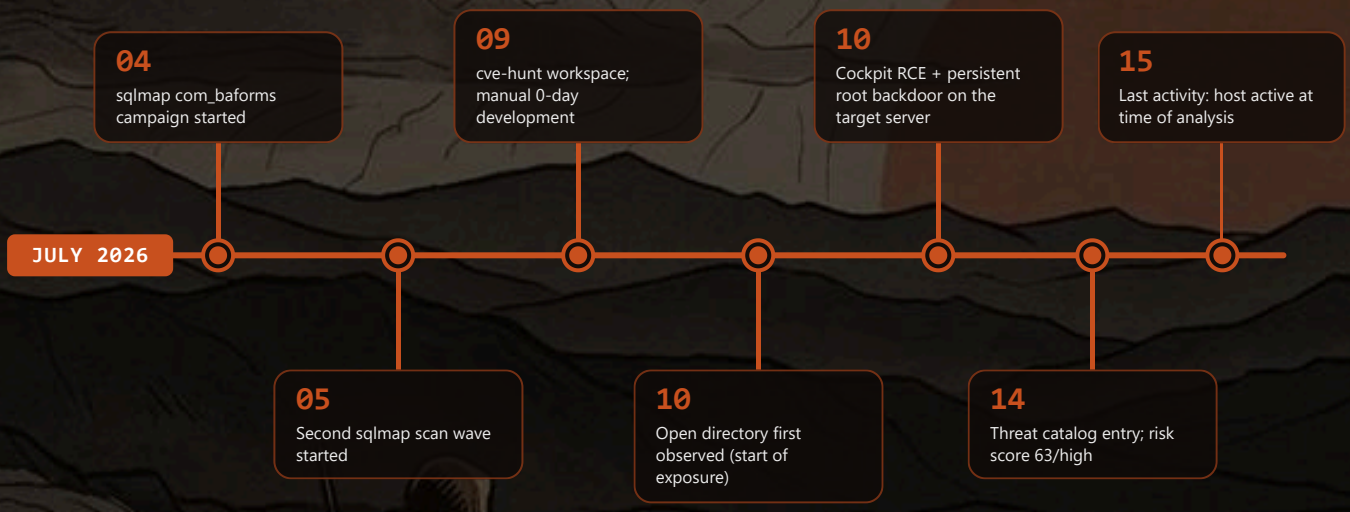
Harvests from GitHub if new CVEs exist (180 s).

4

Targeting Logic

Fresh CVEs that require no authentication and yield **RCE/SQLi** on CMSs are prioritized.

Timeline



SYNTHETIC PHANTOM is one of the notable eCrime actors tracked by etugen.io in 2026 that run the attack cycle largely via an autonomous AI agent. The campaign was executed repeatedly over at least a week.

Part 2: **Campaign Analysis**

Campaign 1 – AI-Automated Mass CMS Exploitation

The actor's primary campaign uses the **shodan-attack-pipeline** skill to find vulnerable WordPress/Joomla/Drupal sites and attempt mass exploitation with sqlmap, nuclei and copied PoCs. The main vector is SQL injection against the Joomla `com_baforms` file-upload endpoint.

> *The autonomous agent scanned a broad target pool (sqlmap_targets 30, manual_targets ~57, joomsport candidate list of hundreds of hosts) with a single operator and no human intervention; the prominent aspect of this campaign is its **automation tempo and scale**.*

The agent continuously updated and prioritized the target pool via the `all_poc_targets.txt` scorecard, performing reconnaissance and triage steps entirely automatically. This is an architecture that multiplies the target breadth reachable by a single operator.

The agent confirmed **2 hosts as "verified vulnerable" for CVE-2026-6433**; these are genuine vulnerability points the actor verified as open to exploitation.

~50
ACTIVELY ATTEMPTED HOSTS

hundreds
SCANNED CANDIDATE TARGETS

2
HOSTS CONFIRMED VULNERABLE

The campaign was executed repeatedly between 2026-07-04 → 07-10 with no human intervention; the autonomous agent scanned the target pool continuously for days and updated itself each round with new CVEs from NVD. This industrial tempo is the distinguishing signature of the AI-orchestrated operation.

Campaign 2 – Targeted Root Compromise and Persistence

The actor **manually obtained root and established persistence** on a server through a targeted chain; this is the report's verified persistent-access gain. The victim is a single, container-based Linux server that runs its own service infrastructure via the Cockpit management panel. The organization's identity has been withheld in the report for victim protection and is hereafter referred to as the **target server**.

Stage	Technique
Initial Access	Cockpit -oProxyCommand injection (:9090), Basic test:test default credentials
Privilege Escalation	Container-to-host escape via nsenter --mount=/proc/1/ns/mnt --pid=/proc/1/ns/pid
Persistence	Actor's public key to /root/.ssh/authorized_keys; reinforced by a cron running every minute
Verification	root SSH to the target server → hostname; whoami; id (success marker 🔥 SHELL:)

The most critical finding: the actor's **passwordless live SSH private key** (ed25519, 432 B) is exposed in the open directory. The key comment is `root@vm168712.hosted-by.qwins.co`; it is both the actor's machine identity and an operational credential.

Because the cron rewrites the key every minute, the standard "delete the key" remediation is ineffective; moreover, since the key has been exposed for months, **anyone other than the actor** can connect to the target server as root; effectively a multi-party compromise.

Campaign 3 – KYC Fraud: Stolen Identity Document Repository

The `/id/` directory on port 7777 contains **identity cards (CNI), passports, residence permits, bank statements and PACS documents** belonging to real French citizens. The documents are combined with an OCR + automated registration pipeline and used in KYC fraud.

1. **Bulk leak pattern:** File names are sequential application numbers (26242, 26779, 26830…) + "emprunteur" (loan applicant); the documents were not leaked individually but **in bulk** from the repository of a French credit/KYC platform.
2. **Automated registration:** bt<number>@mailinator.com accounts + tess_*.PNG OCR artifacts via `register_and_exploit.py` → document verification automation.

> *"Are the passport-style images stolen?" YES. Original documents belonging to dozens of real individuals and accompanying full KYC files (proof of address, bank statement) are present.*

Two Operations, One Workstation

The open directory reveals two parallel criminal operations at once: (A) AI-automated mass web exploitation: broad but inconclusive; (B) KYC fraud with stolen French identity documents: concrete, victim-impacting. Both share the same autonomous agent infrastructure and the same operator identity.

Operation	Status
A: Mass CMS exploitation	Large-scale automated scanning; 2 hosts confirmed vulnerable
B: KYC / identity documents	Concrete stolen PII; tied to automated registration pipeline
Targeted: target server	Verified persistent root breach (manual)

COMMAND & CONTROL INFRASTRUCTURE

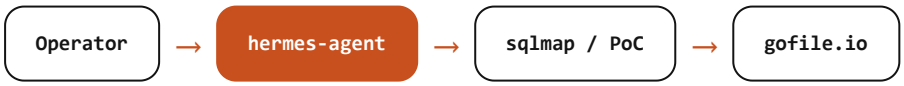
TECHNICAL INFRASTRUCTURE – C2

The actor uses its own autonomous agent infrastructure instead of a commercial C2 framework. Hosting and exfil components:

DETAIL	COMPONENT
`87.58.201.80 (:22, :80, :7777, :9876, :13309)`	Primary workstation
`SimpleHTTP/0.6 Python/3.11.15`	Open directory server (7777, 9876)
`OpenSSH_9.6p1 Ubuntu-3ubuntu13.18`	SSH: direct management
`AS213702 / QWINS LTD`	Hosting (high-risk)
`87-58-201-80-dynamic.dk.customer.tdc.net`	PTR: conflicts with ASN (§verify)
`upload_gofile.py → gofile.io`	Verified exfil channel
`target server:9090 (Cockpit)`	Compromised secondary host

Hosting Contradiction

While the threat record's ASN says QWINS LTD, the PTR record points to a Danish TDC dynamic line. The actor's SSH key comment (`vm168712.hosted-by.qwins.co`) suggests the workspace may have been generated on a QWINS VPS and moved to the machine on the TDC line.



TECHNICAL INFRASTRUCTURE: TOOLING

Most of the arsenal is copied public PoCs; however, the actor's original contribution is **three 0-days** that read and bypass the vendor patch, plus the autonomous agent modules. Actor-produced tools:

TOOL	FUNCTION
hermes-agent (autolearn)	Autonomous AI agent framework + CVE/PoC pipeline
github_search.py (HermesAutolearn/2.0)	GitHub PoC harvesting for a fixed CVE list
sppb_0day.py / pbck_0day.py / yoast_0day.py	Patch-bypassing original 0-days (Russian comments)
salf_backdoor.py	Cockpit RCE + container escape + cron persistence
shell.php / gif_shell.php / minishell.php	12–14 webshell variants (including polyglot)

0-Day Signature Pattern

> All three 0-days carry the same fingerprint: `.htaccess (AddType ... .PHP)` inside a ZIP + upper/lower-case extension bypass (.php .php5 .phtml .PHP .pHp) + `RCE_<id output>_RCE` verification marker → a single author's signature. These patterns make low-false-positive WAF/EDR signatures.

VICTIM TELEMETRY & LEAK

VICTIM ANALYSIS

Victims fall into three categories: only-targeted (not breached) CMS sites, the single verified root breach, and stolen-PII victims.

VICTIM SET	STATUS
Mass CMS targets (~hundreds)	Targeting only, no verified breach
CVE-2026-6433 (2 hosts, IP withheld)	Confirmed vulnerable
Target server (small-scale service)	VERIFIED persistent root breach (Cockpit → host)
French individuals (dozens)	Stolen CNI / passport / bank statement (KYC)

Critical Data Exposed

> The most critical finding: **the actor's own live SSH private key** (salf_key, ed25519) and **genuine French identity documents** in the open directory. Since the key has been exposed for months, the target server compromise is now multi-party; the identity documents, in turn, directly require individual victim notification (TLP:AMBER, PII).

etugen.io

ATTACK CHAIN MAPPING

MITRE ATT&CK MAPPING

The SYNTHETIC PHANTOM attack chain spans from reconnaissance to persistence; MITRE ATT&CK mapping:

- ✕ Recon & Weaponization
- ✕ Initial Access & Execution
- ✕ Escape & Persistence
- ✕ Collection & Exfiltration

MITRE ATT&CK Mapping

TACTIC	TECHNIQUE (ID)	OBSERVED
Reconnaissance	Active Scanning (T1595)	nuclei/Shodan/Netlas mass scanning
Initial Access	Exploit Public-Facing App (T1190)	com_baforms SQLi, Cockpit RCE
Priv. Escalation	Escape to Host (T1611)	nsenter container escape
Persistence	Web Shell / Scheduled Task (T1505.003 / T1053.003)	shell.php + cron SSH key

Analyst note: SYNTHETIC PHANTOM runs the attack cycle with an LLM-based autonomous agent; it automates the reconnaissance–exploitation–exfil chain with the autolearn pipeline. The real impact concentrates in the targeted root access and stolen identity data combined with this automation scale.

INDICATORS OF COMPROMISE

INDICATORS (IOC)

The indicators below have been simplified for safe handling. Verify with your own telemetry before blocking; ASN blocking must be done per-host due to the PTR contradiction.

INDICATOR	TYPE
`87.58.201.80`	Primary workstation
`HermesAutolearn/2.0 · sppb-0day`	User-Agent (agent/0-day)
`RCE__RCE · AddType ... .PHP`	Webshell / 0-day signature
`salf_key (ed25519, exposed privkey)`	Actor credential: MOST CRITICAL

Defense Recommendations

Immediate measures: Block 87.58.201.80 per-host + abuse report to QWINS. Flag the `com_baforms` / `com_sppagebuilder` endpoints and the HermesAutolearn/2.0 User-Agent in the WAF. Clean the crontab on the target server + verify host integrity (key deletion alone is insufficient).

DEFENSE RECOMMENDATIONS

IMMEDIATE MEASURES (0–24H):

- Target server: clean the crontab + audit /root/.ssh/authorized_keys; verify container/host integrity.
- Block 87.58.201.80 per-host; file abuse/notification with QWINS LTD + the target server owner.
- WAF/EDR: flag the HermesAutolearn/2.0 and 'Mozilla/5.0 sppb-0day' User-Agents.
- Distribute the exposed SSH key (salf_key) as an IOC; key rotation on the target server is mandatory.
- Inventory/patch the relevant CMS components (BA Forms, SP Page Builder, PageBuilder CK, Yoast).

MEDIUM & STRATEGIC:

- Hunt for the RCE__RCE wrapper in uploaded PHP files and the in-ZIP '.htaccess AddType ... .PHP' signatures.
- Monitor requests to the com_baforms / com_sppagebuilder SQLi endpoints.
- Match French identity-document victims with the relevant KYC/credit platform and notify (PII).
- Audit Cockpit (:9090) default credentials (test:test) and the -oProxyCommand pattern.
- Proactively scan for open directory (Python http.server) exposure in your own environment.

ETG-2026-002 · **Conclusion**

SYNTHETIC PHANTOM is a financially motivated eCrime actor that automates the attack cycle end to end with an AI agent. The operation's strength lies in the **tempo and scale** a single operator reached through this automation.

The actor's concrete gains concentrate on three axes: **the targeted, persistent root breach on the target server**, **KYC fraud fed by stolen French identity documents belonging to dozens of real individuals**, and **three original 0-days that bypass the vendor patch**. The actor's own OPSEC failure (the open directory) made this rare insight and the exposed key possible.

ASSESSMENT

AI-orchestrated attack automation is now a maturing threat category. In this case the real impact concentrates in **the exposed SSH key, the persistent root on the target server, and the identity-document victims**; defense priority should therefore be these concrete gains.