Hive Pro

Threat Level

Red

Hiveforce Labs

THREAT ADVISORY

VULNERABILITY REPORT

CVE-2026-42208: The LiteLLM Flaw Letting Attackers Reach Deep Inside

Date of Publication	Admiralty Code	TA Number
April 30, 2026	A1	TA2026118

Summary

First Seen: April 2026
Affected Products: BerriAI LiteLLM
Impact: A critical SQL injection flaw in LiteLLM (CVE-2026-42208) is allowing attackers to pull sensitive data straight from backend databases with minimal effort. By abusing a weakness in the API key verification process, threat actors can inject malicious queries through a crafted Authorization header, no authentication required. Observed activity shows targeted attempts to extract high-value assets, including API keys, master keys, provider credentials, and configuration data, using refined techniques like UNION-based queries and schema enumeration. With exploitation already underway and attackers demonstrating a clear understanding of LiteLLM's internal structure, exposed instances should be treated as compromised, making immediate patching and full credential rotation critical.

CVE

CVE	NAME	AFFECTED PRODUCT	ZERO-DAY	CISA KEV	PATCH
CVE-2026-42208	BerriAI LiteLLM SQL Injection Vulnerability	BerriAI LiteLLM	✗	✓	✓

Vulnerability Details

#1

Hackers are actively going after sensitive data stored in the LiteLLM open-source LLM gateway by exploiting a critical flaw, tracked as CVE-2026-42208. This issue falls under SQL injection (CWE-89) and stems from how LiteLLM handles API key verification. During this process, the Bearer token from the Authorization header is directly inserted into a PostgreSQL query instead of being safely parameterized. As a result, an attacker can manipulate the query using a simple quote character and execute arbitrary SQL commands. The flaw impacts versions from 1.81.16 up to 1.83.6.

#2

What makes this vulnerability particularly dangerous is how easy it is to exploit. It does not require authentication or user interaction. An attacker simply needs to send a crafted HTTP request to a LiteLLM proxy endpoint, such as /chat/completions, embedding a malicious payload in the Authorization header. The injection is triggered when the malformed API key fails validation, allowing the payload to execute through the proxy's error-handling logic.

#3

Researchers observed attackers using targeted SQL payloads, including UNION SELECT statements, to extract data from critical database tables. These include LiteLLM_VerificationToken, which stores API keys including the master key, as well as litellm_credentials and litellm_config, which hold upstream provider credentials and environment configuration details. The attacker also employed column enumeration techniques to better understand the database structure, dynamically adjusting payloads to refine access.

#4

The vulnerability was disclosed on April 20, 2026, and exploitation attempts were detected just days later. Activity traced to infrastructure showed a coordinated effort using consistent tooling and a Python-based request pattern. While no confirmed post-exploitation activity has been reported, the level of precision suggests a clear understanding of the data's value. Security researchers strongly advise treating exposed LiteLLM instances as compromised and immediately rotating all API keys, master keys, and stored credentials.

Vulnerability

CVE ID	AFFECTED PRODUCTS	AFFECTED CPE	CWE ID
CVE-2026-42208	BerriAI LiteLLM (>= 1.81.16, < 1.83.7)	cpe:2.3:a:berriai:litellm:*:*:*:*:*:*:*:*	CWE-89

Recommendations

Upgrade LiteLLM Immediately: Update all LiteLLM proxy instances to version 1.83.7 or later without delay. This patched release replaces the vulnerable string interpolation with parameterized queries, eliminating the SQL injection vector in the API key verification path. Organizations should prioritize any internet-facing instances for immediate patching, given confirmed active exploitation within hours of disclosure.

Apply Interim Workaround If Patching Is Delayed: If upgrading is not immediately feasible, set "disable_error_logs: true" under "general_settings" in the LiteLLM configuration. This removes the code path through which unauthenticated input reaches the vulnerable query. Additionally, place the proxy behind a reverse proxy that blocks any Authorization header value containing single quotes, parentheses, or SQL keywords such as UNION, SELECT, FROM, OR, and comment delimiters.

Rotate All Stored Credentials and API Keys: Treat the database of any LiteLLM instance that was internet-reachable while running a vulnerable version as compromised, even in the absence of confirmed extraction. Rotate every virtual API key, master key, and upstream provider credential (OpenAI, Anthropic, AWS Bedrock, etc.) stored in the proxy. Revoke and regenerate all affected credentials at the provider level.

Audit Upstream Provider Billing and Access Logs: Review upstream LLM provider billing dashboards and API access logs for /chat/completions traffic originating from unfamiliar IP addresses in the days surrounding the exposure window. Master-key reuse from an unexpected source IP is the most reliable indicator of credential monetization following a successful extraction.

Restrict Network Exposure of AI Proxy Infrastructure: Move LiteLLM proxies behind an internal network boundary or a mutually authenticated reverse proxy. AI gateways consolidate enough credential value that direct internet exposure is no longer a defensible default. Ensure that port 4000 (or whichever port the proxy listens on) is not publicly accessible without additional access controls.

Potential MITRE ATT&CK TTPs

Tactic	Technique	Sub-technique
Initial Access	T1190: Exploit Public-Facing Application	
Credential Access	T1552: Unsecured Credentials	T1552.001: Credentials In Files
Collection	T1005: Data from Local System	
Discovery	T1082: System Information Discovery	
Resource Development	T1588: Obtain Capabilities	T1588.006: Vulnerabilities

Indicators of Compromise (IOCs)

TYPE	VALUE
IPv4	65[.]111[.]27[.]132, 65[.]111[.]25[.]67

Patch Links

https://github.com/BerriAI/litellm/releases/tag/v1.83.7-stable

https://github.com/BerriAI/litellm/security/advisories/GHSA-r75f-5x8p-qvmc

References

https://github.com/BerriAI/litellm/security/advisories/GHSA-r75f-5x8p-qvmc

https://webflow.sysdig.com/blog/cve-2026-42208-targeted-sql-injection-against-litellms-authentication-path-discovered--hours-following-vulnerability-disclosure

What Next?

At **Hive Pro**, it is our mission to detect the most likely threats to your organization and to help you prevent them from happening.

Book a demo of **HivePro**.

REPORT GENERATED ON

April 30, 2026 • 9:00 AM

More at www.hivepro.com