Hive Pro

Threat Level

Red

Hiveforce Labs

THREAT ADVISORY

VULNERABILITY REPORT

MLflow SSRF Flaw (CVE-2026-64849) Hands Attackers Your Cloud Credentials

Date of Publication	Admiralty Code	TA Number
August 19, 2026	A1	TA2026238

Summary

First Seen: August 17, 2026
Affected Products: MLflow (all versions prior to 3.15.0)
Impact: A critical server-side request forgery (SSRF) vulnerability, tracked as CVE-2026-64849, has been found in MLflow, the widely used open-source platform that machine learning and data engineering teams rely on to track experiments, package code, and deploy models. The flaw lives in an unauthenticated webhook-testing endpoint and affects every version of MLflow before 3.15.0. What makes this especially dangerous is that the server hands the full response back to the attacker, turning a normally "blind" SSRF into a tool that can read sensitive data directly. Security researchers observed real-world attacks against internet-exposed MLflow servers within hours of public disclosure, with attackers racing to harvest cloud credentials and deployment tokens. The issue has been fixed in MLflow 3.15.0.

CVE

CVE	NAME	AFFECTED PRODUCT	ZERO-DAY	CISA KEV	PATCH
CVE-2026-64849	MLflow Server-Side Request Forgery Vulnerability	MLflow (mlflow)	✗	✓	✓

Vulnerability Details

#1 CVE-2026-64849 is a server-side request forgery flaw affecting all MLflow releases before 3.15.0. It sits in the model-registry webhooks feature, specifically the unauthenticated POST /api/2.0/mlflow/webhooks/{id}/test endpoint. A default Tracking Server (mlflow server) runs without authentication on a local SQLite backend, so this webhook API is often exposed to untrusted traffic straight out of the box.

#2 The root cause is an incomplete security check. The _validate_webhook_url() helper added in version 3.10.0 blocks private and cloud metadata addresses, but only inspects the first URL. The delivery code in mlflow/webhooks/delivery.py then follows HTTP redirects and re-resolves the hostname without pinning the validated address. An attacker can host a public URL that passes the initial check and returns an HTTP 302 redirect to a link-local metadata service or internal loopback address. Because the hostname is re-resolved after the check, the flaw is also open to DNS-rebinding attacks.

#3 What turns this from a typical SSRF into a full-read primitive is the test endpoint's behavior: instead of just firing the webhook, it reflects the upstream HTTP status and response body to the caller. The attacker therefore receives the actual contents of whatever internal or cloud resource was reached. On AWS, Azure, and Google Cloud, this exposes temporary IAM credentials, OAuth tokens, and environment configuration from the metadata address at 169.254.169.254, along with internal microservices and loopback admin consoles that trust the host.

#4 Active exploitation is not theoretical. A honeypot network reported attackers targeting internet-exposed MLflow instances within hours of the CVE being assigned. A public proof-of-concept is also available, including a confirmed reproduction against MLflow 3.13.0, and the flaw was independently disclosed in public GitHub issue #24179. The vulnerability is fixed in MLflow 3.15.0, and defenders should assume opportunistic mass exploitation is ongoing.

Vulnerability

CVE ID	AFFECTED PRODUCTS	AFFECTED CPE	CWE ID
CVE-2026-64849	MLflow (Before 3.15.0)	cpe:2.3:a:mlflow:mlflow:*:*:*:*:*:*:*:*	CWE-918

Recommendations

Upgrade MLflow Immediately: Update every MLflow Tracking Server to version 3.15.0 or later without delay. This release fixes the redirect-following and DNS-rebinding weaknesses in the webhook delivery logic and is the only complete remedy for the vulnerability.

Rotate Potentially Exposed Credentials: Because patching does not undo any theft that already occurred, assume exposure on any internet-facing server and rotate all cloud IAM keys, OAuth tokens, and API secrets tied to the host instance. Do this even after you have upgraded, since credentials harvested before the patch remain valid until rotated.

Audit Access Logs for Abuse: Review server access logs for requests directed at /webhooks/ /test, paying special attention to any that predate your upgrade. Requests to this path from unexpected sources are a strong signal that someone attempted or succeeded at exploitation.

Isolate MLflow Behind Authentication: Place Tracking Servers behind an identity-aware proxy or authentication layer and remove them from direct public exposure. Default MLflow deployments ship without authentication, so internal tracking portals should never be reachable from the open internet.

Potential MITRE ATT&CK TTPs

Tactic	Technique	Sub-technique
Reconnaissance	T1595: Active Scanning	T1595.002: Vulnerability Scanning
Initial Access	T1190: Exploit Public-Facing Application	
Credential Access	T1552: Unsecured Credentials	T1552.005: Cloud Instance Metadata API
Resource Development	T1588: Obtain Capabilities	T1588.006: Vulnerabilities

Patch Links

https://github.com/mlflow/mlflow/releases/#release-v3.15.1

https://github.com/mlflow/mlflow/releases/tag/v3.15.0

https://github.com/mlflow/mlflow/pull/24258

References

https://github.com/mlflow/mlflow/security/advisories/GHSA-7gwp-5pfp-969j

https://www.linkedin.com/posts/watchtowr_watchtowr-intel-is-observing-in-the-wild-activity-7495481580723838976-qzUA/

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

August 19, 2026 • 08:20 AM

More at www.hivepro.com